

Why Activists Should Consider Making Lots of Money

By Brian Tomasik

First published: 2006. Last nontrivial update: 2019 Nov 08.

Summary

It's common for activists to equate "nonprofit work" with "doing good" and "for-profit work" with "selling out." But in fact, if you're well suited to working at a corporation or startup company, this option might be better than working at a nonprofit, because the extra money you could donate at the for-profit job would buy new nonprofit employees, while at the nonprofit, you might partially be replacing someone else. That said, I would not advise the "earning to give" approach for everyone, and I think there are cases where people should clearly work on a problem directly rather than making money to donate toward it. There may also be cases in which you don't counterfactually replace another person, in which case the direct impact of the job is potentially significant.

See also "[Employers with Huge Matching-Donations Limits](#)".

Contents

- [Summary](#)
- [Introduction](#)
- [Replaceability?](#)
 - [A model with fixed budgets and wages](#)
 - [A model with variable wages](#)
- [Earning to give](#)
- [Earning keeps options open](#)
- [Cases where earning to give may not be optimal](#)
 - [Differences in productivity](#)
 - [Transactions costs](#)
 - [What's the true opportunity cost of your hires?](#)
 - [Information as the limiting resource](#)
 - [Startup and fringe causes](#)
 - [Risk of becoming affluent](#)
- [Chicken-and-egg questions](#)

- [Don't earn to give if your talent lies elsewhere](#)
- [Career choice is personal](#)
- [General guideline](#)
- [Do what you love?](#)
- [Links](#)
- [Acknowledgements](#)
- [Appendix: My personal history with this topic](#)
 - [My career history](#)
 - [History of the idea](#)
- [Appendix: Reforming from the inside?](#)
- [Appendix: How to supervise others with minimal time](#)
- [Footnotes](#)

Introduction

Activism is about "changing the world" and "making the world a better place"—through research, lobbying, and civic engagement. So it seems quite natural to suggest that someone who wants to effect change should work for an advocacy organization, or a citizen-action campaign, or a charitable foundation. To do something different—for instance, taking a job at a big law firm or corporation—is sometimes seen as "selling out" and abandoning one's moral ideals. This common perception is reflected in the title of one *Wall Street Journal* article: "[Social Principle vs. High Salary: Reconciling Ideals with Financial Goals Can Be a Struggle for Some in Their 20s](#)" (by Adelle Waldmen, 2005).

Replaceability?

However, there is a factor that is often ignored in this line of thinking: Namely, that when an organization hires one person for a competitive position, it thereby chooses not to hire another.

Example. Suppose Alice and ten other people are applying for the position of campaign coordinator for an advocacy group. Alice is hired and feels good about all of the work that she's doing. She forgets, however, that because she was hired, someone else wasn't. In her place, that person would have done all sorts of great work too.

I often have a bias to find my own experiences, thoughts, and projects more important than those of other people because they're mine. They're closer to me and have more sentimental value. I work hard to generate my ideas, while the

millions of other writings on the Internet, activist campaigns, strategy sessions, etc. feel "cheaper" because I didn't put in any effort toward them. But from an impartial standpoint this is not true. Other people's efforts are very important to *them*, and overall impact is determined by the quality/efficacy of the work, not whether I did it myself. When I adopt a broader sense of ownership, in which "we're all in this together" and "your work is my work", I can reduce the sense of favoritism I have toward doing things myself and appreciate the power of earning to pay others to undertake good work.

A model with fixed budgets and wages

One simple model of replaceability is the following. Consider all of the nonprofits that do good work. Those nonprofits have essentially fixed budgets that allow them to hire an essentially fixed number of employees. Since nearly all of these positions will be filled by *someone*, much of the work that Alice does with one of those organizations would have been done by someone else in her place. Hence, the *difference* that Alice makes on the world is only the amount by which she's more productive than whoever would have been hired instead.

This will usually not be quite true. For instance, consider an unrealistic model in which all nonprofits within a certain sector agree on a ranking of all job applicants from best to worst. If the nonprofits have fixed budgets and pay fixed salary amounts for each job slot, they will be able to fill a constant number of slots, and they will choose the highest-ranked applicants to fill those slots. If you then enter the pool of applicants and get hired, and if you aren't the lowest-ranked of the applicants hired, then you don't "knock out" the person who would have been hired in your place, because that person will still be hired somewhere else. Instead, you prevent the lowest-ranked applicant from being hired. In this case, the difference you make is the amount by which you're more productive than that lowest-ranked person who got kicked out would have been. (This ignores what happens to that lowest-ranked candidate; I'm implicitly assuming that he goes off to work for an organization that makes no net impact to the world.)

When we relax the assumptions made in the preceding paragraph, we find that the benefit you make by working for an organization is probably less than the difference between you and the worst employee. For one thing, it won't always be the case that the worst person who would have been hired will be "kicked out"; maybe someone slightly more capable would leave instead. Second, if you're more talented than other candidates, you may receive more remuneration, which lowers the total hiring power of your organization.

A model with variable wages

The assumptions of a fixed charity budget and constant employee wages are unrealistic. In Econ 101, it's taught that when you enter a labor market, you shift the labor supply curve rightward. Then this both decreases wages and increases quantity of labor supplied. If supply is highly elastic, then the increase in quantity supplied that results is rather small. For some charity jobs, such as "grunt" positions that nearly anyone could fill, it's plausible that supply is quite elastic, since people would quickly move elsewhere if wages declined even slightly. Supply is presumably also pretty elastic for corporate jobs, which tend to have very competitively priced wages. In other cases supply might be less elastic, in which case your entering the market suppresses wages enough that the charities or corporations can hire an appreciable fraction of one new worker on balance.

For further discussion of replaceability models, see

- ["Which Ethical Careers Make a Difference?: The Replaceability Issue in the Ethics of Career Choice"](#) by Benjamin Todd
- ["Replaceability"](#) by Paul Christiano
- ["Replaceability arguments are overblown"](#) by Ben Kuhn.

It's important to consider whether replaceability applies to a significant degree with the job you're considering. If not, then the direct impact of the work may be highly significant. For certain industries—including work for military companies, software/hardware companies that help speed up AI, or other companies that [speed up risky technologies](#)—the direct impacts may significantly detract from the good accomplished by your donations.

Earning to give

In some cases, productivity differences between you and your potential replacement could be quite large. In other cases, even small differences in effectiveness could translate into large impacts. If, for instance, Alice were a slightly better fundraiser than the next person who would have been hired, she might be able to bring in millions of extra dollars over her lifetime.

But another option to consider, besides working for a charity, is to work at a corporation or other employer that will provide vastly greater income.

Replaceability suggests that even if the employer is slightly negative in its social impact, not all of this impact will be realized in a counterfactual sense, although

enough will be that you'd presumably still want to avoid the worst industries.

The key benefit of earning to give is that, while your labor would have been somewhat replaced by another person, your donations at a high-paying job probably would not have been replaced, or at least not very much, given that most people don't donate large percentages of their income. With these extra donations, you can create not just a partial new charity job like you would by working for a charity but actually *several* new charity jobs, by expanding charities' budgets.



Following are a few examples of how much nonprofit employees earn in the USA around 2014, taken from Glassdoor.com. (Note: Inclusion on this list does not mean I endorse the charity.) These figures underestimate the actual cost of a nonprofit employee, because they don't count the employer portion of Social Security tax (6.2% of income), Medicare tax (1.45% of income), office/building costs, janitors, human resources, and other maintenance costs that increase slightly when an additional employee is hired.

<i>Organization</i>	<i>Position</i>	<i>Average yearly pay (US\$)</i>
Humane Society of the United States	Communications	\$30K
	Director	\$67K
PETA	Youth Marketing Coordinator	~\$34K
	Administrative Assistant	~\$25K
	Quality Assurance Specialist	\$58K
Brookings Institution	Corporate Counsel - Publications	\$71K
	Research Assistant	\$42K
	Senior Research Analyst	\$81K
	Program Coordinator	\$43K

Oxfam	Campaign Project Manager	\$47K
	Senior Development Manager	\$59K
Human Rights Watch	Coordinator	\$45K
	Researcher	\$69K
	Senior Legal Advisor	\$118K

Earning keeps options open

Another strong case for earning instead of focusing on altruism directly is that wealth is robust against changes in your moral or epistemological outlook. It provides flexibility and keeps your options open. Of course, if there's a risk of your becoming affluent, you may not want your options *too* wide open. One way to navigate this dilemma is by using a donor-advised fund, which offers tax deductions in the short run and requires you to eventually donate your income to a charity of some sort.

In addition, starting out in a corporate career can keep options open relative to starting out in a nonprofit one. Nonprofits would *love* to hire former McKinsey consultants, but McKinsey may be wary of hiring a former animal activist.

Finally, having money can open up doors that wouldn't have been possible without it, including working on your own startup without being squeezed for cash, founding a new nonprofit before you have funding to support it, or even just doing work totally on your own without needing an organization to support you and donors who will hold you accountable to particular expectations.

Cases where earning to give may not be optimal

In practice, I don't think a career of making money to donate is always the obvious choice, for several reasons.

Differences in productivity

Avoiding [Lake-Woebegone bias](#) when assessing your own abilities is important, but it may very well be that you are exceptional in a particular field. Differences in

personal effectiveness can be huge: Some leaders of nonprofit organizations, say, probably are many times more influential than their counterparts. Certain academics are perhaps 3, 5, or 20 times more productive than their colleagues in terms of predictably high-quality work. If, upon reflection, you find that you excel in a certain area that can be used to prevent lots of suffering, then by all means pursue it, whether it makes money or not. (If you are exceptionally skilled at torturing squirrels or cleaning bathrooms with a toothbrush, please do not pursue these vocations.)

Transactions costs

As is true of so many theoretical [Pareto improvements](#), the reality of transactions costs can get in the way. Hiring someone else to do a job for you requires finding that person, establishing trust, maintaining communication to share information, setting up proper incentives to avoid [principal-agent problems](#), etc. Plus, making money means paying taxes twice if your arrangements are sufficiently informal that you can't deduct the payments as donations to a public charity. And if you do set up a 501(c)(3) for tax reasons, that involves yet more paperwork.

Eliezer Yudkowsky [elaborated](#) on this point:

I once observed that Ricardo's Law of Comparative Advantage is the theorem that unemployment should not exist. The gotcha being that if someone is sufficiently unreliable, there is a cost to you to train them, a cost to stand over their shoulders and monitor them, a cost to check their results for accuracy—the existence of unemployment in our world is a combination of transaction costs like taxes, regulatory barriers like minimum wage, and above all, lack of trust. There are a dozen things I would pay someone else to do for me—if I wasn't paying taxes on the transaction, and if I could trust a stranger as much as I trust myself (both in terms of their honesty and of acceptable quality of output). Heck, I'd as soon have some formerly unemployed person walk in and spoon food into my mouth while I kept on typing at the computer—if there were no transaction costs, and I trusted them.

What's the true opportunity cost of your hires?

It can be tempting to say things like the following: "The people at the charity I'm donating to earn \$30K per year. But I can donate \$90K per year. So my donations are supporting 3 new people to do charity work." This may or may not be true

depending on who the people are. For example, say one of the people working for the charity was [Nick Cooney](#). Given his assiduity, acumen, and altruism, if he weren't working for animal organizations, he could and probably would be earning significant money (>\$100K per year?) to donate to animal organizations. Say that if he weren't working for animal organizations, he would donate \$75K to them, but when he does work for them, he takes \$25K in salary. The full cost of hiring Nick Cooney is thus \$100K, not \$25K.

In other cases, the same may not be true. People doing veg outreach or administrative tasks or working for a very popular cause like global health might be hired from a pool of applicants who would have otherwise gone on to do non-altruistic jobs and wouldn't have donated much. Their opportunity costs may be closer to their direct compensation costs.

There are some talented people who could earn a lot in the private sector and wouldn't donate much in that context, but instead they do charity work. Accepting a lower salary is their form of contributing to the cause. I think this situation is fairly common, and maybe it helps explain the halo society associates with working for nonprofits. It's atypical for people to donate substantial fractions of their income; most potential high earners donate their talents instead. Had such people entered the corporate world, they would have been ensnared by the lifestyles of their peers and not focused much on philanthropy. People like this have lower opportunity cost than Nick Cooney, because they wouldn't otherwise be donating a lot, but they might still have some opportunity cost when working for your organization if they would have worked for another good nonprofit had you not hired them.

In general, nonprofit employees have some opportunity cost, even if they wouldn't otherwise be earning much. If your organization's fundraiser wasn't working for you, she might be working for another charity that's almost as good as yours. Your hiring her does two things:

1. Brings her to work for your org instead of where she would have worked counterfactually.
2. Frees up money by the org where she would have worked to hire someone else.

These opportunity-cost calculations can become complex, but take a simplified example: You, Organization X, hire a research analyst, Bob, at \$40K. Had Bob not worked for you, he would have worked for \$30K at Organization Y. Bob has 1.3 times as much impact working for X as if he had worked for Y. Because Y didn't hire Bob, they hire someone else, Chris, who would have done something non-altruistic

altogether had he not been hired. Chris is less productive than Bob would have been, so where Bob would have produced 1.0 of value at Y, Chris produces 0.6. What has the \$40K of pay for Bob bought?

- Had the money not been spent, Bob would have produced 1.0 of value at Y. Chris would have produced nothing.
- With the \$40K being spent, Bob produces 1.3 of value at X, and Chris produces 0.6 at Y.

The net increase in value is $(1.3+0.6)-(1.0+0) = 0.9$. This is less than the "naive" view we would have gotten by looking directly at Bob's producing 1.3 of value in his work at X, though the discrepancy isn't huge, and it obviously depends a lot on the numerical inputs. The main point, though, is that like employees who could be earning and donating a lot, talented nonprofit employees who would otherwise be working at alternate good organizations are more expensive than we might have naively supposed.

Similar considerations apply if your money doesn't fund employees directly but instead is granted toward projects. Naively you might feel as though the dollars you grant toward a project are fully responsible for what that project achieved, when in fact, if you hadn't made the grant, maybe the grantees would have gotten money from someone else (in a similar way as an employee you fund at organization X might instead have gone to work at organization Y).

A charity's project has least opportunity cost if it pays for buying things, like advertising or supplies, that don't take human capital away from other potential lines of work.^a That said, I think most of the important altruism tasks that need doing now (mainly, cause prioritization) involve lots of human capital.

Information as the limiting resource

The argument for making money to donate treats wealth as the limiting resource in efforts to reduce suffering. However, you may decide that the hardest part about making a difference is figuring out where exactly to donate at all. While differences in earnings between those in the nonprofit and private sectors can differ by a multiple of 3, 5, or 15 times, differences in cost-effectiveness of donations can vary by 20, 50, or 100 times. If, by working for a charity, you're able to improve your on-the-ground knowledge about what kinds of activities and groups really are effective, that can be extraordinarily valuable. If, by doing research, you're able to spend time learning about details of an issue in a way that leads you to discover a vastly more cost-effective cause than you ever would have discovered as a lawyer

working on mergers and acquisitions, then you've made a great career choice. For very speculative and under-studied causes, you might be able to spend your whole lifetime researching and still not hit sharply diminishing returns on the value of further information.

Of course, it's worth pointing out that you don't *have to* work on the ground yourself to get advice on what works and what doesn't: You can also make friends with people in that position and consult them. And you don't have to earn a PhD in, say, political science in order to get sound advice about the implications of international policies on prospects for compromise in the future; you can also establish a relationship with people who do know about such topics. While it's important not to assume that every task can be optimally traded in [Ricardian](#) manner, it's also important not to assume that you're the only one who can adequately gather information, assess options, and carry out altruistic projects. Even if information is the limiting resource, it may be best still to make money and hire experts to make decisions for you given knowledge of your personal objective function.

Startup and fringe causes

The case for making money to donate is very strong for established causes where the limiting resource is funding to hire more eager employees—e.g., at [Vegan Outreach](#). The situation is different when it comes to a fringe or startup charity that doesn't yet exist or where it's hard to hire people other than yourself to get the ball rolling, except possibly at high opportunity cost. In these cases, you may have no option but to do the work directly if the value is high enough relative to what you could accomplish by donating to more mainstream causes.

Over the years, I've learned that it's very hard to convince people to do what *you* want them to do; almost always, they do what *they* want to do. The same is true for me as well if someone asks me to do something. This fact of human nature can make it hard to outsource your pet projects, by earmarked donations or otherwise, unless you pay someone a lot of money. Instead, you should generally assume that if you're going to donate toward altruism, you'll probably have to pick which organizations and people you think *already* do good work. It's not easy to convince people to change course and remain passionate while doing so.

Risk of becoming affluent

I have been supposing that the hypothetical donor in this scenario [would live as frugally as possible](#)—purchasing fancy clothes, expensive meals, and other luxuries only to the extent that doing so is justified on the grounds of long-term career

advancement to earn a higher income or to engage in social networking with powerful people. (An investment banker is unlikely to pull in hundreds of thousands per year if she wears a ratty T-shirt to meetings with clients.) Thriftiness can sometimes be hard to maintain amidst social pressure towards luxury. If you find that to be the case—or think, after adjusting for overconfidence bias, that it may become the case in the future—you could consider donating to a foundation or donor-advised fund to enforce your commitment to charitable giving, while simultaneously reaping income-tax deductions in the year the money is earned. Of course, this strategy may not be a good idea if you intend to use the money in ways that can't be supported by a foundation or public charity—e.g., to support your own start-up company before you reach the stage of attracting venture-capital funding.

There may be a risk in encouraging less dedicated altruists to earn to give, since if someone isn't already strongly committed to altruism for the long term, she may stray away from altruism more easily. On the other hand, very committed people tend also to have talents for doing direct altruism work, making it less obvious that they should earn to give.

Chicken-and-egg questions

Direct altruism work may have higher payoff than it appears, especially if it helps to change society's Zeitgeist in the long term. Lots of wealthy philanthropists (and non-donor private citizens) are influenced by the ideas in society at large. By doing outreach, writing, and speaking on important issues now, you might make it so that in a few decades it's not just one philanthropist (i.e., yourself, if you had done earning to give) donating to highly important causes but perhaps many of them.^b The route of social influence is riskier than direct earning, because you might have little impact, or you might (with lower probability) have huge impact. It's not obvious how it compares with direct earning in expected value.

However, there's also the option of earning and in the process funding others to shape society's ideas. You might be able to accomplish more total influencing in the same amount of time this way, so it's possible earning could be better even from the perspective of long-run memetic change.

Or maybe rather than earning yourself, you should bring others into the movement, some of whom will earn. Or you could write about building movements to encourage people to earn. Or you could earn to fund people to write about building such movements. And so on. We run into issues of whether the chicken or egg comes first here.

This becomes confusing to think about, but it helps to remember that each of these strategies takes time and has bounds on how far it can go before other limiting factors kick in. Direct outreach can be limited by funding. Earning is bounded by how many people you can fund. Movement-building is limited by funding and having some content to share (otherwise, what are you building a movement around?). Producing high-quality research can attract talented people who stumble upon it. Direct altruism work is constrained by funding and knowledge of which causes are best to support.

There's nothing magic about any of these options. All are a priori plausible possibilities, and the choice depends on what factors are limiting right now, what you're best at, and what you'd most enjoy.

Don't earn to give if your talent lies elsewhere

The perspective of earning to give largely treats altruism projects as given + scalable vs. money as the limiting resource. A different view could argue for unique abilities of individuals as the limiting resource, with talent attracting money when it's warranted. I think each of these extremes isn't quite true, but there's some wisdom to both. I do think that really impressive projects will tend to attract funding, and there are some extraordinary amounts of wealth in the world to be courted. If you have a special knack for an altruistic undertaking—research, outreach, fundraising, etc.—I suspect you'll be able in the long run to acquire the finances you need. At least, this has been my experience.

There may be exceptions in cases where few other major donors share your values, and if so, you might need to fund the project yourself. However, I'm doubtful whether this is often the case; somewhere in the world there must be wealthy people who are on your wavelength, and if not, maybe your particular project doesn't have high odds of success anyway. (Thanks to Jonah Sinick for this point.) In most cases, I think extraordinary people should focus on the area of altruism in which they're extraordinary and leave the funding situation to work itself out later.

In the case of earning to give, you can, by donating more than your counterfactual replacement, expand the total size of the philanthropy pie. If you rely on being exceptional at something else, you're mainly displacing existing philanthropy money (unless you do such novel work that a philanthropist decides to donate more in total). This is somewhat a concern but maybe not a big one, depending on the counterfactual value of where the philanthropist would have given instead.

Because one altruist who earns to give can support many altruists doing direct work, in general, we should see most altruists doing something other than earning, unless we assume that the supply of non-earning altruists is already far higher than the number that can be employed. Holden Karnofsky [echoed](#) this sentiment:

I have a general intuition that over the long run, [effective altruists] EAs are going to be disproportionately well funded and will be more bottlenecked by human than financial capital (this doesn't mean the best thing to do is work directly on EA activities, but it would argue against earning to give all else equal, at least if you don't expect to be one of the top earners [...]).

As of 2015, I'm uncertain if I agree that EAs will be well funded in the long run. Right now most of the organizations I'm involved with or support are mostly money-limited, though a few other EA organizations are talent-limited. The question depends partly on how skilled the work you want to fund is and how perfectionist you are about hiring only the best people to do that work. I think it's fair to say that most charities are usually money-limited.

Career choice is personal

Whether earning to give is a good idea depends not just on your talents but on your personality and interests. For instance, Simon K [explains](#) that he stopped earning to give mainly because

1. he found it hard to remain passionate about high-level business work that didn't have much connection to altruism
2. he wasn't sure there was enough room for funding at organizations that shared his values
3. it was hard for him to socialize with colleagues who didn't share his values.

These are great points to consider, although it's also important not to assume you wouldn't enjoy earning to give prematurely. If you're uncertain, try earning first, since it's usually easier to switch from business to activism than the reverse direction.

When I was earning to give at Microsoft, I didn't find Simon's points particularly troublesome:

1. In my personal case, I don't lack intrinsic motivation when doing non-altruistic work. One of my original reasons for thinking about earning to give was that in

2005, I assumed that the best direct altruism work I could do would be to become a political activist, and that seemed less intellectually stimulating than doing geeky work at a corporation. As of 2015, I think I enjoy direct altruism work about the same as earning to give. I can get excited about almost any topic as long as I know that it's useful, and software in particular is really interesting to me. (Coding can be one of the most enjoyable experiences in life.)

2. As mentioned in the previous section, the charities I want to fund appear relatively money-constrained. This is particularly true for organizations like the [Humane Slaughter Association](#).
3. I didn't find value differences from colleagues to be a concern, and in fact, debating philosophy with others at Microsoft tended to solidify friendships, since I went from being "some other coworker" to "a coworker with whom I have interesting discussions". The mileage probably varies depending on the company culture, but in the tech world, being idiosyncratic often isn't a problem.

General guideline

- If you think there are lots of effective charities to fund and not enough money for them, then you should either earn or persuade other altruists to do more earning/donating.
- If you think there's lots of money and not enough effective charities to attract that money, then you should work directly on those charities, improving them to the point where they are worth funding and then soliciting donations.
- If you think money and uses of that money are roughly in balance, then choose whether to earn vs. work directly depending on your comparative advantage, i.e., are you more or less suited (in terms of talent, interest, etc.) for earning relative to direct work than other altruists?

Do what you love?

When I was young, my mom kept a notebook of interesting observations I made. Somewhere around age ~5 I said something to this effect: "It's good that different people enjoy different things, because then different people can do different jobs."

There's a healthy debate over whether and to what extent "do what you love" is good career advice.

- Penelope Trunk [calls](#) this principle "absurd", noting that we love many things and learn new tastes. Instead she recommends to "do what you are" based on your strengths.
- William MacAskill [says](#) that "'Follow your passion' is the stupidest career advice I've ever heard" because many people are passionate about unproductive activities. Instead, he recommends doing something valuable, which may ultimately make you happier anyway.

That said, there are real differences in how much different people enjoy different careers. And as Vince Broady [notes](#): "If you don't personally care about what you are doing, you are not going to be competitive at it."

In general, less enjoyable careers should pay more (at least holding intelligence, education, and social status constant; working at a slaughterhouse pays terribly despite being one of the worst jobs I can imagine). The fallacy of [generalizing from one example](#) may incline us to assume that if we don't enjoy something, then no one does. But this isn't necessarily true. While aggressive, non-stop, real-time trading on Wall Street sounds very stressful to me, one trader I talked with said he found the atmosphere exciting, and I've heard other people complain that their finance jobs weren't stressful enough. While I find big-picture research more interesting than concrete, administrative work, other people feel the opposite way. Some people live and breathe the stock market or the latest tech gadgets, while others live and breathe physics or philosophy. My own preferences have even changed over time.

In some cases, earning to give can be a way to do what you love. Prior to 2005, I had assumed that I would have an obligation to work for an activist organization in some capacity, even though I loved math more. After realizing that making money to donate could be at least as good, I was happy to know that I might not have to choose between making a difference and enjoying my work.

On the other hand, many of the highest-earning careers are not very enjoyable, which may present a conflict between doing what you love and doing the most good.

Personally I think that doing what you love, within bounds of reason, is a strong heuristic. Unless you're one of a small fraction of people who can stick to a rational plan on the basis of willpower and goal-directed focus, you're likely to burn out if you push yourself too hard. When you do what you love, you may be several times more productive than someone else who has little enthusiasm for the same work. That said, like most heuristics, doing what you love may not be the best advice in all

cases. And if you're lucky, what you love (or what you can come to love over time) may encompass a very broad range of activities.

Links

- [80,000 Hours](#)
- [Choosing a Conventional Career to Help Animals](#) by Patrick Brinich-Langlois
- [Peter Singer Discusses the Ethics of Shopping](#)
- [Money: The Unit of Caring](#) by Eliezer Yudkowsky
- [Utilitarian careers](#) discussion on Felicifia
- [Salary or startup? How do-gooders can gain more from risky careers](#) by Carl Shulman

Acknowledgements

My understanding of this topic has been refined by the many people with whom I've discussed career choice, both inside the effective-altruism movement and beyond.

Appendix: My personal history with this topic

My career history

When I was in high school (2001-2005), I was a member of a club called Youth Ending Hunger. One of our activities was to raise money, through selling lollipops, bagels, and pizza. Pizza sales were the most lucrative, typically bringing in ~\$150 per sale. To hold one, we had to (1) request approval for a table in the hallway after school, (2) tell everyone in the club the date and times of the sale, (3) order pizza, (4) request permission to leave class 15 minutes early to set up, and (5) serve pizza at the table for about an hour. We needed 3-5 people to help out. All told, this probably implied at least ~5-7 hours of people time, not including the accounting paperwork later. Meanwhile, I had a friend whose father probably earned ~\$100 per hour at his job. I often thought to myself, "Why don't we ask him for a donation of 1.5 hours of his time and skip the whole pizza sale?"

In Dec. 2004, I began thinking about replaceability in career choice, when I heard of a friend who had quit working for a corporation because the friend considered it "evil." Upon reading Peter Singer's "[Famine, Affluence, and Morality](#)" in summer

2005, I also realized the high potential value of money. This was reinforced when I came across the following [observation](#) from Jack Norris around the same time:

Many activists view money and wealth as evil. As a source of power, money can be used to promote either evil or good. Just think how much better the animals would be if vegans had significant amounts of money. If each vegan had enough money to buy and distribute multiple copies of educational materials, the animals would greatly benefit. Someone who works a job that isn't directly promoting animal rights, but who can use their money to fund the resources needed by our movement, will be doing much to help the animals.

In fall 2005 I began to research high-earning careers for myself. In Dec. 2005, I wrote a document: "[Choosing a Utilitarian Career](#)" (.doc). In 2006, I wrote a shorter, more pedantic piece called "[Utilitarian Career Choice](#)" (.tex).

Throughout college I explored many possible high-earning careers. I realized that a corporate career path would be seen as odd by my do-gooding activist friends, so I often explained my reasoning to them. In fall 2007, I wrote an editorial for my college newspaper titled "[Making Money and Making Change Not Mutually Exclusive](#)."

Eventually I chose to pursue a career in software. One reason was that it offered the potential for pursuing a startup with very high expected earnings later on. Another motivation was that it had high starting salaries (often higher than entry-level finance/consulting remuneration, though with a slower growth trajectory) and didn't require extra education (in contrast to law/medical school). I realized that I might not earn indefinitely—that I might want to earn for a few years and then "retire" to do independent writing—which is why higher starting compensation was important.

I ended up working at Microsoft's Bing search engine for four years. I enjoyed the work a lot, but gradually it became apparent that I should return to altruism research full-time, because I had too many topics I needed to study in order to figure out how to do good with donations. Another reason was I was told repeatedly that my writing and thinking were of sufficiently high-quality that I should consider doing it directly as my area of comparative advantage.

History of the idea

Jeff Kaufman [cites](#) me as one of the first people to write about earning to give on the web, and while it's true that I came up with the argument on my own, it seems

that *many* people have developed the same idea independently. I later discovered at least three other friends who had already pursued the same path for similar reasons. As Eliezer Yudkowsky [notes](#), the basic principle is not original:

There *is* this very, very old puzzle/observation in economics about the lawyer who spends an hour volunteering at the soup kitchen, instead of working an extra hour and donating the money to hire someone to work for five hours at the soup kitchen.

Earning to give is basically a combination of the concept of comparative advantage with the observation that jobs in an economy will tend to be at least somewhat replaceable.

In 2011, the organization [80,000 Hours](#) began popularizing the idea of earning to give to a wider audience. Actually, at that point, they called it "professional philanthropy." However, in Aug. 2012, Jeff Kaufman expressed his worry that the name was confusing and asked for alternatives. I suggested "earning to give," and the rest of the community liked it, so the name stuck. Niel Bowerman independently came up with a similar phrase: "earn to donate." The name has since become well known through a [Washington Post piece](#) and many other citations, and now [has its own Wikipedia article](#).

Appendix: Reforming from the inside?

Sometimes altruists propose going into seemingly non-altruistic careers with the intent of reforming an institution from the inside. I haven't studied in detail whether this can be effective, but I have a weak intuition that it's probably not optimal, unless you also plan to go into the same career for other reasons (e.g., making money). My guess is that working on the issue directly from the outside is often better. Following are some possible arguments against reforming from the inside:

1. *Reaching an influential position in a big company or government agency is hard.* Many people want those jobs for other reasons. Of course, influencing from the outside is hard too, but my point is that going the inside route isn't obviously an easier path to success. If you are skilled enough to rise in the ranks internally, chances are you can also be very successful externally.
2. *Constraints on internal reform.* Even if you do reach an influential position, it's still not clear you can reform too dramatically. A company CEO is fundamentally constrained by what's profitable. A politician is fundamentally constrained by donors and popular opinion (unless it's her last term in office).

When you lobby from the outside, you can do whatever you want. There's a famous quote by a politician (which I can't find at the moment) in which he tells activists something to the effect of "Thank you for making me enact this policy." In other words, he wanted to do it, but he couldn't until the political winds changed such that doing so became politically feasible.

3. *Too much front-loaded investment.* It takes years or decades to become influential in an industry or agency, and in that time, your ideas about what cause to focus on may very well have shifted. You also need to spend lots of time on activities not relevant to your goal.
4. *Less public conversation.* Along the same lines, a significant part of the value of advocacy work comes from sparking conversation and refining our ideas about what we think the right policies are. That debate is easiest to hold in the public domain. If you speak too openly about your intentions from inside the organization, you may not be promoted in that organization. (Of course, there are some exceptions, like organizations that actively seek new ideas about how they can grow.) By building the movement from outside, you can more directly focus on the aspects of the situation that you care about and can generate momentum that may eventually influence those already in the organization to change, either voluntarily or by public pressure.
5. *Less honest.* I feel like if you choose to work for organization XYZ, you should generally act in the interests of organization XYZ while employed with them. Going into organization XYZ with the secret intention of shifting its priorities in more than small ways is plausibly morally dubious, relative to common-sense moral intuitions. (I would make exceptions for things like whistleblowing. I think taking undercover footage of cruelty on factory farms falls into that category.)
6. *Less precedent?* Most of the examples of reform that come to my mind have involved campaigning from the outside. Of course, this availability heuristic may be flawed, perhaps because the outside campaigns are more salient, more interesting to historians, or just more common in general. Certainly plenty of bureaucrats, politicians, and business executives have aimed to make many internal reforms. That happens all the time when new leadership comes to power. But typically this is done by people who intended to build their careers in that field for other reasons. Maybe politicians are the best exception: Some people deliberately run for office to challenge the status quo from the inside. This seems like a decent approach if you're suited to politics and want to work on issues that are politically feasible. (Insect suffering, for instance, is not a politically feasible issue in the near future.)

7. *Risk of losing one's ideals.* If you go work with a group of people who have different moral priorities than you, are you certain that you'll be the one to influence them rather than them influencing you? Maybe social pressure will cause you to become less hard-core about your previously held moral values. (Of course, sometimes losing one's idealism due to a better understanding of the unpleasant realities of how the world works can be a good thing, if it makes you more effective at actually accomplishing your goals or makes your goals more reasonable.)

All of that said, there may indeed be cases where working on the inside is the best approach, such as if you have a lot of detailed technical contributions you want to make, if you want to learn from an established group of professionals before striking out on your own, or if there are other factors that change the calculations considerably.

A long time ago I read a biography of Ralph Nader. It mentioned an episode in which Nader met a young man who was planning to reform a corporation from the inside by working for it. Nader encouraged him to lobby from the outside instead, which the man did.

In 2015–2016, [a wave of companies](#) pledged to go cage-free due to pressure from a handful of activists at a few animal organizations: "nearly 200 U.S. companies – including every major grocery and fast-food chain – that together buy half of the 7 billion eggs laid monthly have pledged to use only cage-free eggs by 2025." It seems very hard to imagine that these activists could have achieved the same scale of victory by becoming executives at McDonald's, Walmart, etc. and then pushing for internal reform. Part of the reason is that if the pressure is only internal, then other executives and shareholders can raise the complaint that the reforms hurt the bottom line. In contrast, if the pressure comes externally, then *not* doing humane reforms becomes worse for the bottom line.

Appendix: How to supervise others with minimal time

Suppose you're earning to give in order to fund research on important topics. Since your job may occupy most of your time, it can be tough to also oversee the research work you're paying for. Following are some possible suggestions for nudging the research in useful directions without requiring too much oversight.

1. Create a master list of topics you want explored. (An example is the Foundational Research Institute's "[Open Research Questions](#)".)
2. Create a reading list of material with which researchers should be familiar.

3. If you yourself have written on the topics of interest, encourage your researchers to read your past writings so that they can get up to speed. (For instance, basically all of my original thoughts about altruism appear on my websites and hence—at least in theory—don't need to be conveyed in person.)
4. Find methods for researchers to vet the quality of their output. An informal way is to post writings on forums with voting functionality and comments. A more formal approach is to submit the essays to academic journals or popular media websites; doing this [takes more time but also](#) has higher payoff. Writing [on Wikipedia](#) is another good option if the topic is important.
5. Convert the papers that your researchers produce [into audio files](#) so that you can listen to them in your down time (e.g., when walking, shopping, preparing food).
6. Remember that micromanagement is often not productive. People are more passionate when exploring topics they personally care about.
7. Join the organization's board of directors, and provide high-level feedback at board meetings.

Footnotes

- a. The opportunity cost of stealing human capital away from other good organizations is a consideration that only altruists worry about. If you selfishly only care about your organization, then this opportunity cost doesn't matter, except possibly indirectly for game-theoretic reasons. [\(back\)](#)
- b. One prominent example of this is the impact of Nick Bostrom's *Superintelligence* book, which—in conjunction with other previous groundwork that had been laid—led to a \$10 million donation from Elon Musk [announced in Jan. 2015](#) and then another £10 million from the Leverhulme Trust [announced in Dec. 2015](#). [\(back\)](#)