

1	Introduction	1
1.1	The importance of the question	2
1.2	Intuitions for waiting	3
2	Direct efforts to increase near-term human welfare	4
2.1	Public projects and the Ramsey formula:	
	“Discount by $\delta + \eta g$ ”	4
2.1.1	On the relationship between the interest and social discount rates	5
2.1.2	Empirical estimates	5
2.2	The argument from the EA community:	
	“Discount by $\eta g^* > \delta + \eta g$ ”	5
2.3	A proposed correction:	
	“Discount by $\eta g < \delta + \eta g$ ”	7
2.4	Complications	8
3	Model with two actors and constant effectiveness	10
3.1	Motivation	10
3.2	Results	11
3.3	Some numbers	16
4	Model with one actor and varying effectiveness	18
4.1	Hingyness	18
4.2	Basic model	19
4.3	Markov-process model	21
4.3.1	With arbitrary η_P	21
4.3.2	With $\eta_P = 1$	23
4.3.3	Some numbers	23
5	Further considerations	24
5.1	Relationships to endogenize	24
5.1.1	Learning	24
5.1.2	The relationships between r , g , and fund size	24
5.1.3	The relationship between δ_P and fund size	25
5.1.4	Uncertainty about long-run r , δ_P , and η_P	25
5.1.5	Bargaining over the future	25
5.2	Optimal investment	26
5.2.1	Mission hedging	26
5.2.2	Movement building	27
5.2.3	Accounting for exogenous movement growth	27
6	Conclusion	28
7	References	28

Discounting for Patient Philanthropists

Philip Trammell

August 18, 2019

IN PROGRESS – PLEASE DO NOT CITE

Abstract

Philanthropists must decide to what extent to spend their resources on present philanthropic projects and to what extent to invest them, at the available interest rate, for use on future philanthropic projects. I investigate some features of this discounting problem in the case that the philanthropist is patient—i.e. has a low (including zero) rate of pure time preference. In particular, I explore the problem in some detail in the case that the philanthropist will spend his resources on direct efforts to improve near-term human welfare. I also take some steps toward outlining the problem in the case that the philanthropist will spend his resources on effecting a long-term “social trajectory change”, including for example by mitigating risks of human extinction. In both cases, standard economic assumptions imply that patient philanthropists should invest most of their resources in most circumstances.

1 Introduction

Most people act roughly so as to maximize their own future welfare, discounted according to some positive and substantial “rate of pure time preference” (Olson and Bailey, 1981).¹ They will tend to avoid pain where possible, for instance, but will accept severe pain in the future to avert moderate pain the present. We will call someone a “philanthropist” if his goal is to maximize discounted total welfare, not only his own.² We will call a philanthropist “patient” if he has low (including zero) pure time preference with respect to the welfare of his philanthropy’s beneficiaries. Thirty-eight percent of economists who study discounting (Drupp et al., 2018) and most moral philosophers (Broome, 1994) agree that, at least when setting policy and engaging in philanthropy, we ought to be patient.

The Effective Altruism (EA) movement brands itself in large part as a community of patient philanthropists. There is currently no consensus within the movement regarding whether philanthropic funds are best spent, from a patient perspective, on direct efforts to improve others’ wellbeing in the

¹Find more/better references for this

²We will thus assume for simplicity that philanthropists are discounted total utilitarians. Many of the observations presented here are in fact more widely applicable.

near term or on more speculative efforts to improve the distant future.³ Nevertheless, the consensus in both camps is that we should give now, and indeed that the question is so settled that further investigation is of low priority. For efforts focused on the near term, it is typically argued that, once we consider the rate at which the world’s poor are getting richer or, relatedly, the flow-through effects of relieving their poverty, sooner gifts do more good than later gifts, even at compound interest. For efforts focused on the long term, it is typically argued that we are living at a special time in history, during which our efforts to influence the future are uniquely impactful (see e.g. Parfit, 2011), and that the specialness of our time is fading more quickly than the interest rate. Some appear recently to be rethinking this position (see e.g. Vivaldi, 2019), but advocating long-term investment is still far from mainstream in the EA community, and very little effort has been put into its implementation.

Most small donors do in fact spend their charity budgets as soon as they earn them (CITE). Perhaps ironically, however, most large philanthropists outside the EA movement, who typically have made no explicit commitment to patience, already spend as slowly as the law permits—in the United States, that is 5% per year—and this behavior has historically met with strong criticism from those in the EA movement (see e.g. Karnofsky, 2007).

My aim here is twofold. First, I hope to argue emphatically that the question of when to give is important and underexplored. Second, I hope to argue tentatively that patient philanthropists should collectively invest most of their funds for giving later, regardless of whether patient philanthropic resources are best spent on near-term-focused efforts or on efforts to affect the distant future.

1.1 The importance of the question

If our aim as philanthropists is to maximize global welfare, absent pure time preference, what should we do with our money? At first glance, this is an intimidatingly complex question. Given the wide variety of giving opportunities, we must split the space of options many times. We must first decide on a moral framework within which to compare causes; we must then find the cause which promises to produce the most non-discounted moral value; we must then find the interventions that most effectively further this cause; and we must finally find (or found) and fund the organizations most efficiently implementing these interventions.

When we begin by asking whether to “give now or later”, however, the question simplifies. It is not, I think, obvious whether the answer is more likely *a priori* to be *now* or *later*. Asking whether to wait thus splits the space of options roughly evenly, in the way that asking whether to “give to Charity X or elsewhere”, for any particular X, does not. And if we conclude that the answer is *later*, we have almost solved the problem of global prioritization.

Not fully. Even if we conclude that we ought to give later, we must still decide where to invest in the meantime. I discuss this question briefly in §5.2. But except in the knife-edge case that giving now and giving later are about equally valuable, the proportionate difference in expected philanthropic impact achievable across reasonable investment strategies (e.g. investing everything in

³Cite EA Survey.

the US total market) is not nearly as great as the proportionate difference in expected philanthropic impact between any reasonable approach to giving now and any reasonable approach to giving later. The timing question is thus extremely decision-relevant, far out of proportion to the attention that patient philanthropists have yet devoted to it.

1.2 Intuitions for waiting

A strong case for exploring the possibility of investing to give later would hold even if, as proposed above, the case for giving later were neither obviously weaker nor obviously stronger than the case for giving now. But in fact there are compelling informal arguments for thinking that, at least in ordinary circumstances, patient philanthropists should prefer waiting. Here are three.

First: The future contains many moments, and the present contains only one. Just as the worthiest beneficiary of our charity is probably not the beggar we happen to be passing, therefore, the worthiest charitable projects are probably not those available at the moment in which we happen to find ourselves. Just as those without a preference for local philanthropy are willing to wire their funds a long distance, we should expect philanthropists without pure time preference to be willing to invest their funds for a long time.

Second: In the future, we will all be dead. Patient philanthropists thus have much stronger preferences for what happens to the distant future than self-interested actors (or impatient philanthropists) do—stronger, that is, than those of any other actors alive today. Conceiving of “the future” as a large, four-dimensional object, we should expect that if it were for sale, patient philanthropists would be willing to outbid everyone else for it. But in fact it is for sale: we can, for instance, buy land and sell a hundred-year lease on the land, the latter of which typically costs almost as much as the former (Giglio et al., 2015). In other words, we can acquire the right to determine what is done with land from 2119 until the oceans dry, almost for free. Of course, longer leases sell at even less of a discount. More generally, we can give the impatient our assets in the short term in exchange for a larger share of their assets in the long term. This is what it is to invest.

Third: Patience cannot remain unusual forever. The patient, of course, will invest more than the impatient, and as they do, they will come to own more of the world. Once most of the world belongs the patient, the opportunity to buy the future from its current owners will be gone. The giving opportunities available to us today may be more or less than typically impactful, but we know at least that we are living at a special time with respect to the value of saving.

We will now explore the results of some simple models designed to weigh the above considerations for and against giving later, and other considerations like them, in the context of near-term efforts and the context of longtermist efforts respectively.

2 Direct efforts to increase near-term human welfare

2.1 Public projects and the Ramsey formula:

“Discount by $\delta + \eta g$ ”

Though the question of discounting for patient philanthropists appears to have received little academic attention, there is an analogous question about which volumes have been written. Policy-makers face the question of when it is worthwhile to tax people, thereby lowering their present consumption, to fund projects that increase their future consumption. In response, economists have developed an extensive literature aimed at answering when this is worthwhile—i.e. what “social discount rate” governments should use in trading off consumption at different periods. This literature centers around what is known as the Ramsey formula (Ramsey, 1928).

The Ramsey formula holds that public projects’ consumption payoffs should be discounted at a roughly constant rate $\delta + \eta g$ per year, where g is the rate of annual per-capita consumption growth, η parametrizes the extent to which marginal utility diminishes in consumption, and δ is some rate of pure time preference. It follows from the following conditions:

1. Individuals are homogeneous, and their welfare is roughly “isoelastic” in consumption. That is, it can be represented by some function $u(c) = \frac{c^{1-\eta}}{1-\eta}$ for some η .⁴
2. Consumption growth proceeds at some roughly constant annual rate g .
3. Individuals act so as to maximize their expected future welfare, discounted at some roughly constant annual rate δ , and their policymakers should do the same.⁵

Given Condition 1, $u'(c) = c^{-\eta}$. This implies that, if consumption is multiplied by x , marginal utility in consumption is multiplied by $x^{-\eta}$. Given Condition 2, it holds that, in t years, consumption is multiplied by e^{gt} . It follows from these two observations that, in t years, marginal utility in consumption is multiplied by $e^{-\eta gt}$. In other words, marginal utility in consumption falls at an instantaneous rate of ηg per year. Or: in the absence of pure time preference, individuals would be indifferent between a marginal unit of consumption and $e^{\eta gt}$ units at time t , since each produces equal non-discounted welfare. Finally, given Condition 3, individuals with pure time preference are indifferent between a marginal unit of consumption now and $(e^{\delta t})(e^{\eta gt}) = e^{(\delta + \eta g)t}$ units of consumption at time t . In other words, they discount consumption at the instantaneous annual rate

⁴Isoelastic utility in the case of $\eta = 1$ is defined, via a limit condition, as $u(c) = \ln(c)$.

The assumption of isoelastic utility is made largely for reasons of mathematical tractability, but it does have at least some empirical justification. As Kaldor (1957) famously observed, interest rates and economic growth rates have been roughly constant throughout the developed world since the Industrial Revolution. These “stylized facts” have roughly held to this day. If individuals’ time preferences are time-separable roughly exponential (as they must be if they are stationary and time-consistent), then utility in consumption must be roughly isoelastic, at least over the range of consumption-levels spanned over the last two centuries. This follows directly from the Ramsey formula.

⁵Note that, under these conditions, η simultaneously equals (a) individuals’ coefficients of relative risk aversion, (b) individuals’ elasticities of intertemporal substitution, and (c) the degree of social inequality aversion in consumption.

$\rho = \delta + \eta g$. If policymakers are to discount consumption as their constituents do, they must discount by ρ as well.

2.1.1 On the relationship between the interest and social discount rates

As defined above, ρ is the rate at which individuals discount consumption. Observe that, under ordinary circumstances, ρ will equal the market interest rate r . At least when evaluating projects whose payoffs' time-profiles are comparable to those of some private investment, therefore, a policymaker aiming to discount like her constituents can generally discount by r as observed directly, instead of trying to estimate the inputs to the Ramsey formula.

A policymaker more patient than her constituents—i.e. one discounting welfare at some rate $\delta_P < \delta$ —must be able to separate δ from the other components of the observed interest rate so as to maximize consumption discounted by $\delta_P + \eta g$. The optimal patient policy, however, is perhaps not immediately obvious. A naive “patient policymaker” might be inclined to fund any project whose stream of consumption-payoffs, discounted by $\delta_P + \eta g$, exceeded its cost. But as long as there is a market for private investments offering returns at rate r , there is no sense in funding projects that offer less. The planner should just fund public projects down to the point that their returns equal the market interest rate; then she should split public expenditures between funding lower-return public projects and subsidizing private investment (as explained e.g. by Barrage, 2018), so that the marginal returns to both kinds of projects is always the same.

Furthermore, if the patient government does not have the option to subsidize private investment (or tax private consumption), it must acknowledge that the resources taxed or borrowed to fund public projects will displace present consumption only at the taxpayer's marginal propensity to consume m ; fraction $1 - m$ of the resources used for a lower-return public project will instead displace private investments that would have earned the market interest rate. The patient planner should therefore fund public projects down to a rate of return of $m\delta_P + (1 - m)\delta + \eta g$, which will lie between her own discount rate and that of the impatient citizens. Bradford (1975) explores this problem and works out what effective rate to use in more general circumstances.

The problem a patient philanthropist faces is somewhat different. As we will see below, however, he too must consider the possibility of displacement along similar lines.

2.1.2 Empirical estimates

[For the rest of §2, I will assume values of $\delta = 0.02$, $g = 0.02$, and $\eta = 1.2$. CITE.]

2.2 The argument from the EA community:

“Discount by $\eta g^* > \delta + \eta g$ ”

As noted above, it is common in the EA community to argue that, once we consider the rapid rate at which the world's poor are getting richer or the flow-through effects of relieving their poverty, we should conclude that sooner gifts are preferable. Let us consider these arguments in turn.

The “vanishing opportunities” argument is that the world’s poorest are getting richer so quickly that, in effect, the cost of a unit of human welfare in the developing world is rising more than the interest rate.⁶ Letting g^* denote the consumption growth rate among the world’s poorest, this amounts roughly to the claim that $\eta g^* > \delta + \eta g$.⁷ In other words, the faster-than-average growth in the developing world is claimed to trump the extent to which investor impatience boosts interest rates. Giving now is thus found to do more good than giving later will do, at least before incorporating considerations about learning, value drift, and so on.

Furthermore, even after incorporating other considerations, bloggers in the EA sphere typically find the one-year discount rate to be higher than the one-year market interest rate (see e.g. Harris, 2018). Dickens (2019) recently examined some of the considerations as well, and concluded unusually that the answer was at least unclear—i.e. that we might expect to do better by giving later—but only by making what seem to me like a number of important mistakes in the process, which tend to bias in favor of giving later.⁸ In short, if we are deciding to give to poverty causes now or in a few years’ time, an important and plausibly dominating consideration is that it is quickly getting more expensive to help the global poor.

The “flow-through effects” argument is made perhaps most directly by Giving What We Can, an EA-affiliated organization that encourages people to give at least 10% of their incomes to effective charities. “The money you give has long term benefits to the communities which receive it”, they say. “Curing someone of disease today means that they will be able to contribute much more to their society and its economy from now on, increasing the effectiveness of your donation. In effect, donating earlier does more good more quickly than donating later” (Giving What We Can, 2013).

⁶Alexander (2013) reports Elie Hassenfeld, a cofounder of EA charity evaluator GiveWell, making essentially this point.

⁷Note that if this is the case, then the poor will not be indifferent to trading off present and future consumption at the market interest rate, regardless of how little pure time preference they exhibit, but will actively seek to borrow. The “vanishing opportunities” argument thus implicitly relies on the argument that the poor face credit constraints which prevent them from borrowing at the market interest rate.

Note also that this implies not only that we should give as soon as we can, but that we should borrow as much as we can, if our goal is to help the global poor.

⁸To give one example: he writes “An exponential discount rate p relies on the core assumption that increasing someone’s income by a certain percentage does the same amount of good no matter their absolute income level—if someone’s twice as rich, you need to give them twice as much money to produce the same increase in welfare. In other words, utility is logarithmic with income.” In fact, an exponential discount rate just relies on the assumption that utility is isoelastic. (Recall that logarithmic utility is the special case of isoelastic utility with $\eta = 1$.) Estimates of η vary widely, but most are at least somewhat higher than 1, and it is common to use estimates as high as 2. (See Footnote 2 below for a brief survey of the literature.) In comparing marginal utilities between the global rich and the global poor, choice of η makes a massive difference. For instance, if $\eta = 1$, money goes 100 times as far for someone making \$400 per year as for someone making \$40,000 per year; if $\eta = 1.2$, the multiplier is 250.

This is not the place for a thorough response to Dickens’s work. I just want to cite it as an example of existing thought on this question, and to say that I think it comes to the right conclusion even if for the wrong reason.

2.3 A proposed correction:

“Discount by $\eta g < \delta + \eta g$ ”

Let us grant, for the sake of argument, that $\eta g^* > \delta + \eta g$ as claimed. It follows from this that we would do more good for the poor by giving this year than by giving next.

We cannot, however, naively extrapolate this one-year relationship. When considering whether to give now or in *many* years, what matters is our guess about the long-run relationship between the rate at which our resources will grow (or shrink) and the rate at which doing good with a unit of resources will grow more (or less) costly.

Over the course of a long enough horizon, the rate of increase in the cost of producing a unit of welfare as efficiently as possible cannot, on average, exceed ηg . Otherwise, the most efficient way to produce a unit of welfare would eventually be more costly than one particular way to do so: namely, by giving money to ordinary investors for their own consumption. And since the long-run average rate of increase in the cost of welfare is bounded above by ηg , investing at $\delta + \eta g$ must eventually result in an endowment able to buy more welfare than the endowment we started with. In other words, in the long run, we should discount by ηg , which will be less than the market interest rate as long as market participants, on average, exhibit positive pure time preference.

For illustration: Suppose a dollar can currently buy 1 unit of welfare for someone in the developing world, and suppose consumption in the developed world ($\sim \$40,000$ per year) is roughly 100 times consumption in the developing world ($\sim \$400$ per year). If $\eta = 1.2$,⁹ then a unit of welfare is roughly 250 times more expensive in the developed world; a dollar can buy ~ 0.004 as many units of welfare here as there. (This is consistent with estimates cited in MacAskill (2015) and Ord (2017).) If $\delta = 0.02$, and if this impatience rate can be expected to influence interest rates for a sufficiently long time, then our dollar can buy more than one unit of welfare simply by giving it to developed-world investors after $\log_{1.02}(\frac{1}{0.004}) \approx 279$ years.

This holds even if consumption growth persists at the current rate of roughly 2% per year. In this case developed-world investors will after 279 years be roughly $1.02^{279} \approx 250$ times richer, consuming an average of roughly \$10 million per year. Nevertheless, the corresponding compound interest at $0.02 + 1.2 \cdot 0.02 = 4.4\%$ per year will multiply each dollar invested today by a factor of

⁹Recent research, focusing on the risks people are willing to take in the context of the labor market, implies that in general, coefficients of relative risk aversion are roughly 1 (Chetty, 2017). Global survey data on the relationship between consumption and self-reported wellbeing suggests the same (Gandelman and Murillo, 2014). A survey of 173 domain-expert economists (Drupp et al., 2018) produced a mean estimate of 1.35 and a median estimate of 1.

Notably, however, financial data on the extent to which people are willing to accept safer investments despite lower returns are often interpreted to support the conclusion that people are much more risk-averse, with η -values from 2 to as high as 5. To give a sense of how extreme these estimates are: if $\eta = 3$, marginal consumption for someone consuming \$30,000 per year produces one thousand times more welfare than marginal consumption for someone consuming \$300,000 per year. In my estimation, this anomalous financial data is better explained by the hypothesis (defended by e.g. Weitzman (2007)) that market participants assign substantial probability to a “negative tail” of catastrophes bigger than any historically observed, which threaten to render risky investments valueless. In the face of this risk, individuals demand high returns on risky investments, along the lines of what they would demand if they had higher coefficients of relative risk aversion.

$1.044^{279} = 165,000$, and this effect swamps the satiation effect, as we have seen.

More realistically, furthermore, observe that the rates of return available to those in the developing world, and therefore the rate of increase in the cost of doing good in the developing world, can only be higher than the world interest rate because of some sort of market failure. In a world of efficient and frictionless markets, the world's poor would borrow (and perhaps locally invest) until their consumption growth rate satisfied $r = \delta + \eta g^*$. In a sufficiently poor region, the transaction costs of lending might well be high enough to prevent this equilibration. Indeed, solving this purported market failure is one of the primary motivations for microfinance (CITE). Though implementations of microfinance have generally proven less efficient than many had hoped in the early 2000s (CITE), we can infer that, if a part of the developing world is in fact growing significantly more quickly than the world at large, it will relatively soon be able to enter the world credit market, and the rate of increase in the cost of doing good there will subsequently fall below the world interest rate.

2.4 Complications

We have so far made a number of simplifying assumptions in order to highlight the way in which the compounding impatience of typical investors renders investment appealing from a patient philanthropic perspective. Two of these are necessary and, I believe, potentially problematic: (1) that a positive gap between the market interest rate and the rate of growth in the cost of welfare for investors—that is, investors' impatience—will persist, and (2) that the supply of investment opportunities is and will remain elastic enough that, to a first approximation, additional investment increases total investment. The second assumption is correct in a single-factor “AK” economy; in a more realistic economy with multiple, complementary factors of production (say, labor and capital), increasing the capital stock will increase output at a diminishing rate. Furthermore, part of the increased output will be paid as wages, rather than interest, because in the presence of more capital the marginal productivity of labor will increase. The first assumption appears even more troublesome. The impatience incorporated into interest rates appears to have been decreasing over time (CITE), and it should be expected to fall steadily, as an ever larger proportion of the world's resources are accumulated by more patient actors. (Note that, on its own, a patient philanthropic fund would be unlikely to affect the world interest rate substantially over the relevant time-scale. See the brief discussion of Assumption 5 below.) Weakening these assumptions would push back the date at which giving to the future rich could be expected to create more welfare than giving to the present poor. [Work out more precisely how to think about this, and come up with some rough estimate of the extent to which this changes the story.]

We have also made four other assumptions which may appear crucial but which in fact, it seems, are not. These are (3) that there is no uncertainty, (4) that exponential consumption growth will persist, (5) that the impact achieved by spending at a given time does not diminish with the size of the expenditure, and (6) that gifts will be “consumed” when they are given, without altering recipients' investment decisions. As we will see, scrutinizing and weakening these assumptions does not overturn the central point that patient philanthropists in an impatient world should invest rather

than give.

[Regarding uncertainty, point out formally that: The argument for waiting does not rely on there being no uncertainty, and in fact the presence of uncertainty can substantially strengthen the case. To see this, observe that uncertainty will, to a first approximation, affect everyone similarly. It will therefore be compensated in investment returns. For example, to the extent that there is some annual risk of human extinction, this will push up all returns. To the extent that there is some specific risk of, say, a permanent stock market collapse, this will push up stock returns. We may therefore come to the conclusion that we are very likely to have lost all our money after 279 years, but only along with the conclusion that, conditional on *not* losing our money, it will have grown by enough to compensate us for the risk. Furthermore, this symmetry only holds precisely with respect to structural uncertainty about long-run interest rates, growth rates, etc., as discussed by Gollier (2016). Uncertainty in the magnitude of short-term fluctuations in returns will frighten the impatient more than the patient, and will give the patient an opportunity to earn even *higher* expected long-run returns by choosing high-beta investments.]

Likewise, the argument for waiting does not rely on the assumption that exponential consumption growth will persist. It only requires the two conditions stipulated above: that the positive *gap* between the market interest rate and the rate of growth in the cost of welfare for investors will persist, and that the supply of investment opportunities at some interest rate $r > \eta g$ is and will remain at least somewhat elastic. To illustrate this, consider a scenario that is in some sense the opposite of persistent exponential growth: suppose that consumption peaks in 2100 at a global average of \$50,000 per person, at which point we will reach “technological maturity”, and the only task remaining to civilization will be to consume the remainder of the earth’s resources. If people continue to discount welfare at 2% per year, then they will still consume their resources too quickly. At an interest rate of 0%, individuals will spread their consumption over time such that marginal welfare in consumption increases at 2% per year; consumption will thus decline at roughly $g = \frac{-\delta}{\eta} = -1.7\%$ per year. A patient philanthropist will then be able to create 2% more welfare every year he waits before giving. Hoarding resources at 0% interest will be an appealing philanthropic investment. Finally, as will by now be clear, analogous reasoning applies to any other potential growth path. If consumption growth is cyclical, for instance, following booms and recessions, it may be tempting for a philanthropist to act countercyclically, waiting in times of prosperity and spending in times of poverty. But as long as individuals remain impatient during the hard times, a patient philanthropist will still do better to invest at the prevailing interest rate. The shape of the growth path has no effect on this conclusion.¹⁰

Next, we have been assuming that the welfare the philanthropist can create by increasing investors’ consumption grows linearly with the size of the philanthropic expenditure. This is a reasonable assumption when the expenditure is only a small fraction of total world consumption, since

¹⁰Patient *policymakers* may do best to act countercyclically. By doing so, they may be able to speed growth during recessions; if so, they can then recoup the cost of their interventions via higher tax revenues in the future (CITE). For a patient philanthropist, however, the fruits of a temporary boost to growth achieved by some countercyclical intervention cannot be recouped. They belong forever after to the impatient public.

additional funds can provide marginal consumption increases to an ever wider pool of beneficiaries. (Recall that in this scenario, we are not sorting beneficiaries by poverty.) Invested funds will grow more quickly than world consumption by rate $r - g$ per year, so given enough time, this marginality assumption will not hold. It would likely hold, however, over the relevant time-scale. Total world consumption is currently roughly \$70 trillion (CITE); after 279 years of compounding 2.4% more quickly than total consumption, even \$1 billion would come to equal only about 0.94% of total consumption.

Finally, we have been assuming that gifts at a given time produce immediate consumption increases at that time. In fact, of course, individuals would consume only some fraction m of a marginal transfer and would invest the remainder. Relatedly, we have been assuming that gifts are unanticipated. In the event that they are anticipated, a patient philanthropist can generally do no more good by giving to impatient investors later than by giving sooner, since the impatient beneficiaries will simply invest one dollar less for themselves and their inheritors for each dollar the philanthropist invests on their behalf. If the amount invested on their behalf exceeds the total they would have invested on themselves, then, as long as markets are complete, beneficiaries will invest nothing on their own behalf and borrow against the anticipated future transfer. The patient philanthropist, however, can do almost as much good (or, if $\eta < 1$, even more good) as he can through unanticipated giving by using his funds to subsidize global investment over time. The model below outlines how this could work.

[I intend to lay out all these assumptions more systematically, so that I can be sure (and so that it can be clear) that I have responded to any potential objection.]

The argument from flow-through effects fails for similar reasons. [Here I want to say something like: “As Christiano (2013) observes, the rate γ at which a philanthropic intervention increases consumption growth cannot, in the long run, exceed g . If it did, the flow-through effects of the intervention would eventually be larger than the rest of the economy, and historical acts of philanthropy would have pushed up the growth rate to γ . But the interest rate r is higher than the growth rate, and has been as far back as records go (CITE). So investing to give later will eventually do more than giving now and trusting in the flow-through effects.” But this point needs to be spelled out more precisely.]

3 Model with two actors and constant effectiveness

3.1 Motivation

As noted above, one argument against investing for the sake of future spending on a cause is that, when one’s future spending is anticipated, it can displace investment by of the cause’s less patient funders. This is perhaps most straightforwardly illustrated with the cause of increasing human consumption—beneficiaries might invest less on their own behalf—but the logic applies much more generally. Consider for example the case in which NASA, our “beneficiary”, wishes to spend its

asteroid deflection budget so as to minimize expected asteroid damages subject to some high discount rate, and a patient philanthropist wishes to minimize expected asteroid damages subject to a lower discount rate.¹¹

If the philanthropist can commit to a particular policy, or can subsidize future spending by the beneficiary in an arbitrary (not necessarily linear) way, then, given full knowledge of the beneficiary’s budget and discount rate, the philanthropist can do at least as well as to implement a non-distortionary subsidy schedule: that is, a schedule that allows him to direct his own funds to the times at which he wants to spend them, without altering the beneficiary’s spending schedule (EXPLAIN / CITE). To test the strength of the “substitution” argument against long-term investment, therefore, let us examine a simple model in which, at any moment, the philanthropist can only (a) spend resources, (b) invest resources at an exogenous interest rate r , (c) transfer resources to the beneficiary, (d) linearly subsidize spending by the beneficiary, or (e) linearly subsidize investment by the beneficiary. This model will endogenously account for complications (5) and (6) above. It will not allow the cost of doing good to vary exogenously with time.

As we will see, even in this setting, the philanthropist can adopt a non-distortionary, dynamically consistent policy—that is, one not requiring a commitment mechanism—so long as impact is logarithmic in spending. Furthermore, following this policy can allow the philanthropist to multiply his impact by several *orders of magnitude* beyond what he can achieve by spending naively. (The beginnings of an attempt to generalize this policy to the case in which impact is an arbitrary isoelastic function of spending can also be found below.)

3.2 Results

Let us denote the beneficiary’s discount rate δ , the beneficiary’s starting budget B , the philanthropist’s discount rate $\delta_P < \delta$, and the philanthropist’s starting budget B_P .

Proposition 3.1 (*The optimal individual spending schedule*)

Suppose an individual has isoelastic utility in consumption parameterized by η , a constant discount rate δ , and a budget B , and suppose he can invest his resources at a constant interest rate r . Then the individual maximizes discounted utility by following consumption schedule

$$x(t) = B \left(\frac{r\eta - r + \delta}{\eta} \right) e^{\frac{r-\delta}{\eta}t}.$$

Proof: Let $y(t)$ denote the resources allocated at time 0 for investment until, followed by consumption at, t . Since utility in consumption is time-additive, differentiable, and concave, resource allocation y will maximize utility iff, for some constant k ,

$$\frac{\partial}{\partial[y(t)]} \left[e^{-\delta t} \frac{(e^{rt}y(t))^{1-\eta}}{1-\eta} \right] = k \quad \forall t. \quad (1)$$

¹¹I will assume that this is positive, but it can represent factors such as exogenous annual extinction risk, or the annual risk that the philanthropist’s funds are expropriated, rather than pure time preference.

Taking the derivative and rearranging, we have

$$y(t) = k^{\frac{-1}{\eta}} e^{\frac{r-r\eta-\delta}{\eta}t}. \quad (2)$$

Subjecting this resource allocation to the budget constraint, we have

$$\int_0^{\infty} k^{\frac{-1}{\eta}} e^{\frac{r-r\eta-\delta}{\eta}t} dt = B; \quad (3)$$

$$k = \left(B \left(\frac{r\eta - r + \delta}{\eta} \right) \right)^{-\eta}. \quad (4)$$

Substituting (4) into (2), and observing that $x(t) = e^{rt}y(t)$, we have

$$x(t) = B \left(\frac{r\eta - r + \delta}{\eta} \right) e^{\frac{r-\delta}{\eta}t} \quad (5)$$

as desired. ■

Proposition 3.2 (*The optimal price schedule when taxation is feasible*)

If the philanthropist were able to tax consumption in addition to subsidizing it, he would maximize δ_P -discounted utility by setting the effective price schedule

$$\hat{f}(t) = e^{(\delta_P - \delta)t} \left(\frac{B}{B + B_P} \right) \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right).$$

Proof: When the beneficiary faces some price schedule $f(t)$, he allocates his budget according to allocation $y_f(t)$ such that, for some constant k ,

$$\frac{\partial}{\partial[y_f(t)]} \left[e^{-\delta t} \frac{(e^{rt}y_f(t))^{1-\eta}}{1-\eta} \right] = k \quad \forall t. \quad (6)$$

Taking the derivative, substituting \hat{f} for f , and rearranging, we have

$$y_{\hat{f}}(t) = k^{\frac{-1}{\eta}} e^{\frac{r-r\eta+\delta_P\eta-\delta\eta-\delta_P}{\eta}t} \left(\frac{B}{B + B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}}. \quad (7)$$

Subjecting this resource allocation to the beneficiary's budget constraint, we have

$$k^{\frac{-1}{\eta}} \left(\frac{B}{B + B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}} \int_0^{\infty} e^{\frac{r-r\eta+\delta_P\eta-\delta\eta-\delta_P}{\eta}t} dt = B; \quad (8)$$

$$k^{\frac{-1}{\eta}} \left(\frac{B}{B + B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}} = B \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta} \right). \quad (9)$$

Substituting (9) into (7) and remembering that $x(t) = e^{rt}y(t)$, we see that the beneficiary now spends according to schedule

$$x_{\hat{f}}(t) = B \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta} \right) e^{\frac{r+\delta_P\eta-\delta\eta-\delta_P}{\eta}t}. \quad (10)$$

At each t , the beneficiary thus consumes

$$\frac{x_{\hat{f}}(t)}{\hat{f}(t)} = (B + B_P) \left(\frac{r\eta - r + \delta_P}{\eta} \right) e^{\frac{r-\delta_P}{\eta}t}. \quad (11)$$

From Proposition 3.1, this is the consumption schedule which maximizes the δ_P -discounted utility achievable with total initial endowment $B_P + B$. It follows both that the given price schedule \hat{f} is feasible and there is no superior feasible price schedule. ■

The above may be useful advice for a patient policymaker. Of course, however, a philanthropist cannot set $f(t) > 1$; a philanthropist cannot tax. If B_P is large enough that $\hat{f}(0) \leq 1$, this constraint is never binding. (Observe that $\hat{f}(t)$ is decreasing in t .) Otherwise, let us consider the beneficiary's response if the philanthropist invests all her resources until $t^* : \hat{f}(t^*) = 1$, setting $f(t) = 1$ for $t < t^*$ and $f(t) = \hat{f}(t)$ for $t \geq t^*$.

Let B_t and $B_{P,t}$ denote the beneficiary's and the philanthropist's endowments at t , respectively.

Proposition 3.3 (*The impatient response to the patient philanthropist*)

Suppose the philanthropist chooses $f(t) = 1$ —i.e. invests all her resources—so long as

$$B_{P,t}/B_t < \eta(\delta - \delta_P)/(r\eta - r + \delta_P),$$

and subsidizes at prices given by $f(t) = \hat{f}(t)$ for $t \geq t^*$, where t^* is the time at which the term above reaches equality. The beneficiary will then respond by spending according to schedule

$$x_f(t) = \begin{cases} DB \left(\frac{r-r\eta+\delta}{\eta} \right) e^{\frac{r-\delta}{\eta}t} & t < t^* \\ B \left(\frac{r\eta-r-\delta_P\eta+\delta\eta+\delta_P}{\eta} \right) \left(\eta \frac{B}{B_P} \frac{\delta-\delta_P}{r\eta-r+\delta_P} \right)^{\frac{r\eta}{D(r\eta-r+\delta)}-1} e^{\frac{r+\delta_P\eta-\delta\eta-\delta_P}{\eta}(t-t^*)} & t \geq t^* \end{cases},$$

and subsidization will begin at

$$t^* = \frac{\eta}{r\eta - r + \delta} \ln \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right) / D,$$

with

$$D = \frac{\ln \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right)}{W \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \frac{r - r\eta + \delta}{r\eta - r - \delta_P\eta + \delta\eta + \delta_P} \ln \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right) \right)},$$

where $W(\cdot)$ denotes the product logarithm.

Proof: Suppose that there is some t^* such that $f(t) = 1 \forall t \leq t^*$. If the beneficiary is spreading his consumption so as to maximize δ -discounted utility, he will be indifferent between allocating marginal consumption to $t = 0$ and to any $t \in [0, t^*]$. That is, as in (5),

$$x_f(0)^{-\eta} = e^{(r-\delta)t} (e^{rt} x_f(t))^{-\eta}; \quad (12)$$

$$x_f(t) = x_f(0)e^{\frac{r-\delta}{\eta}t}. \quad (13)$$

Let us now define

$$D \triangleq \frac{x_f(0)\eta}{B(r - r\eta + \delta)}, \quad (14)$$

so that

$$x_f(t) = DB\left(\frac{r - r\eta + \delta}{\eta}\right)e^{\frac{r-\delta}{\eta}t}, \quad t \leq t^*. \quad (15)$$

For the beneficiary, the marginal utility to allocating consumption to any time $t \leq t^*$ is now the same as that of increasing consumption at $t = 0$, namely

$$\left(DB\left(\frac{r - r\eta + \delta}{\eta}\right)\right)^{-\eta}. \quad (16)$$

Now, observe that if the beneficiary's endowment is being spent at rate $D\left(\frac{r\eta - r + \delta}{\eta}\right)$ while otherwise growing at rate r . At time t , therefore, his endowment is of size

$$B_t = Be^{(r - D\left(\frac{r\eta - r + \delta}{\eta}\right))t}. \quad (17)$$

As stipulated in the statement of the proposition, the philanthropist will therefore invest her entire endowment until

$$t^* : \frac{B_P e^{rt^*}}{Be^{(r - D\left(\frac{r\eta - r + \delta}{\eta}\right))t^*}} = \frac{\eta(\delta - \delta_P)}{r\eta - r + \delta_P}; \quad (18)$$

$$t^* = \frac{1}{D} \ln\left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right) \frac{\eta}{r\eta - r + \delta}. \quad (19)$$

Substituting (19) into (15), we have the beneficiary's spending rate at t^* :

$$x_f(t^*) = DB\left(\frac{r - r\eta + \delta}{\eta}\right) \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r-\delta}{D(r\eta - r + \delta)}}. \quad (20)$$

Substituting (19) into (17), we also have

$$B_{t^*} = B\left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1}. \quad (21)$$

Let us substitute (20) for B into (10) and so re-index t^* to 0. Since the philanthropist begins subsidizing to achieve price schedule \hat{f} for $t \geq t^*$, we have

$$x_f(t) = B\left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta}\right) \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1} e^{\frac{r + \delta_P\eta - \delta\eta - \delta_P}{\eta}(t - t^*)}, \quad t \geq t^*. \quad (22)$$

We now have a second expression for the beneficiary's spending rate at t^* :

$$x_f(t^*) = B\left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta}\right) \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1}. \quad (23)$$

Setting (20) equal to (23) and solving for D , we have

$$D = \frac{\ln\left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)}{W\left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \frac{r - r\eta + \delta}{r\eta - r - \delta_P\eta + \delta\eta + \delta_P} \ln\left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)\right)}. \quad (24)$$

Terms (15), (22), (19), and (24) constitute our result. ■

Proposition 3.4 (*An optimal policy when $\eta = 1$*)

If $\eta = 1$, the philanthropist maximizes δ_P -discounted utility subject to budget constraint B_P by investing until t^* and then subsidizing according to price-schedule \hat{f} (as given by Propositions 3.3 and 3.2, respectively). That is, letting f^* denote the optimal price schedule,

$$f^*(t) = \begin{cases} 1 & t < t^* \\ \hat{f}(t) & t \geq t^* \end{cases}.$$

Proof: Given logarithmic utility, the beneficiary's marginal utility in allocating resources toward consumption at any time is independent of f :

$$\frac{\partial}{\partial y(t)} \left[e^{-\delta t} \ln \left(\frac{e^{rt} y(t)}{f(t)} \right) \right] = \frac{\partial}{\partial y(t)} \left[e^{-\delta t} \ln(e^{rt} y(t)) \right] = \frac{e^{-\delta t}}{y(t)}. \quad (25)$$

The beneficiary's consumption allocation $y(t)$, and spending schedule $x(t)$, are thus also independent of f . The philanthropist can therefore set the price schedule so as to allocate her endowment directly to augmenting consumption at the time-periods of her choosing, without worrying about substitution on the part of the beneficiary.

When the philanthropist follows price schedule

$$\hat{f}(t) = e^{(\delta_P - \delta)t} \frac{B_{t^*}}{B_{t^*} + B_{P,t^*}} \frac{\delta}{\delta_P}, \quad t \geq t^* = \ln \left(\frac{B(\delta - \delta_P)}{B_P \delta_P} \right) / \delta, \quad (26)$$

the marginal δ_P -discounted utility in increasing consumption at $t \geq t^*$ is

$$\frac{\partial}{\partial [y(t)/\hat{f}(t)]} \left[e^{-\delta_P t} \ln \left(\frac{e^{rt} y(t)}{\hat{f}(t)} \right) \right] = \frac{1}{(B_{t^*} + B_{P,t^*}) \delta_P}, \quad (27)$$

which is independent of t . And given $f(t) = 1 \forall t < t^*$, the marginal δ_P -discounted utility in increasing the allocation for consumption at $t < t^*$ is certainly less than that in increasing the allocation for consumption at $t \geq t^*$. This follows from the fact that the marginal utilities would be equal if the philanthropist were able at t to impose the schedule of taxes and subsidies described in Proposition 3.2, which would decrease consumption at t and increase consumption at all $t \geq t^*$. (Remember that, by (25), neither the taxes nor the subsidies would be distortionary.)

Since utility in consumption is time-additive and concave, this completes our result. ■

(I believe that following the price schedule defined in Proposition 3.3 is the optimal dynamically consistent policy for arbitrary η , but I have not yet been able to demonstrate this.)

Proposition 3.5 (*Payoff ratio under the optimal policy when $\eta = 1$*)

If $\eta = 1$, the ratio between philanthropist's payoff from following price schedule f^* and his payoff from an immediate transfer to the beneficiary (or an anticipated future gift against which the beneficiary

can borrow) approaches infinity as the ratio between the philanthropist's starting endowment and the beneficiary's starting endowment approaches zero.

Proof: The payoff from following f^* is

$$\begin{aligned} & \int_{t^*}^{\infty} e^{-\delta_P t} \left(\ln \left(\frac{1}{f^*(t)} B \delta^{(r-\delta)t} \right) - \ln \left(B \delta^{(r-\delta)t} \right) \right) dt \\ &= \int_{t^*}^{\infty} e^{-\delta_P t} \ln \left(\frac{1}{f^*(t)} \right) dt \\ &= \frac{\delta - \delta_P}{\delta_P^2} \left(\frac{B_P}{B} \frac{\delta_P}{\delta - \delta_P} \right)^{\frac{\delta_P}{\delta}}. \end{aligned} \quad (28)$$

The payoff from an immediate transfer is

$$\begin{aligned} & \int_0^{\infty} e^{-\delta_P t} \left(\ln \left((B_P + B) \delta e^{(r-\delta)t} \right) - \ln \left(B \delta e^{(r-\delta)t} \right) \right) dt \\ &= \int_0^{\infty} e^{-\delta_P t} \ln \left(\frac{B_P}{B} + 1 \right) dt \\ &= \frac{\ln \left(\frac{B_P}{B} + 1 \right)}{\delta_P}. \end{aligned} \quad (29)$$

Dividing (28) by (29) (and twice applying L'Hôpital's Rule), we have

$$\lim_{\frac{B_P}{B} \rightarrow 0} \frac{\frac{\delta - \delta_P}{\delta_P} \left(\frac{B_P}{B} \frac{\delta_P}{\delta - \delta_P} \right)^{\frac{\delta_P}{\delta}}}{\ln \left(\frac{B_P}{B} + 1 \right)} = \infty, \quad (30)$$

as desired. ■

3.3 Some numbers

Let us first consider the case of a philanthropist aiming to do good by increasing global human consumption-based utility.

Suppose that $\eta = 1$, $r = 0.04$, $B = \$300\text{T}$ (approximately total world wealth (Credit Suisse Research Institute, 2018)), $B_P = \$20\text{B}$, $\delta = 1.5\%$, and $\delta_P = 0.5\%$ (the fully patient rate if the half-life of a fund designed for the long term is about 200 years). Then, by Propositions 3.3 and 3.4, the philanthropist maximizes δ_P -discounted utility by investing until year $t^* = 687$ and then subsidizing consumption according to price-schedule \hat{f} . Doing so will produce consumption increases after year 687 of

$$\begin{aligned} & \left(\frac{1}{\hat{f}(t-687)} - 1 \right) x(t) \\ &= \left(e^{(\delta - \delta_P)(t-687)} \frac{(B_{687} + B_{P,687}) \delta_P}{B_{687} \delta} - 1 \right) B_{687} \delta e^{(r-\delta)(t-687)} \\ &= \left(e^{0.01(t-687)} \frac{(300\text{T} \cdot e^{0.025 \cdot 687} + 20\text{B} \cdot e^{0.04 \cdot 687}) 0.005}{300\text{T} \cdot e^{0.025 \cdot 687} \cdot 0.015} - 1 \right) 300\text{T} \cdot e^{0.025 \cdot 687} \cdot 0.015 e^{0.025(t-687)} \end{aligned} \quad (31)$$

$$= 1.29 \cdot 10^{20} (e^{0.035(t-687)} - e^{0.025(t-687)})$$

for an expected δ_P -discounted utility payoff of

$$\int_{687}^{\infty} e^{-0.005t} \left(\ln \left(1.29 \cdot 10^{20} \cdot e^{0.035(t-687)} \right) - \ln \left(1.29 \cdot 10^{20} e^{0.025(t-687)} \right) \right) dt = 12.89. \quad (32)$$

By contrast, giving the philanthropist's entire endowment immediately (or equivalently, if the gift is anticipated and markets are complete, naively) allows beneficiaries to consume it as would be optimal given pure time preference of δ . This produces an payoff of

$$\int_0^{\infty} e^{-0.005t} (\ln((300T + 20B) \cdot 0.015e^{0.025t}) - \ln(300T \cdot 0.015e^{0.025t})) dt = 0.013, \quad (33)$$

for a payoff ratio of $12.89/0.013 = 992$.

Note that we are comparing (a) a policy aimed at increasing long-term global consumption to (b) an untargeted increase in world wealth today, rather than to (b') a transfer to the world's poorest today. As 992 is substantially higher than the above-cited estimates of the extent to which money does more good in the developing world, however, these figures suggest that patience ("getting the time right") is worthwhile even if it comes at the cost of all targeting ("getting the place right").

Note also that the 687-year waiting period is substantially longer than the fund's "half-life". In other words, the highest-expected-utility policy is one whose impact is most likely zero.

Finally, it might seem absurd to assume that the interest rate or the expropriation rate will remain constant over such a long timespan. As explained by Gollier and Weitzman (2010), however, introducing uncertainty about the long-run values of these variables generally *strengthens* the case for investing. Intuitively, this is because, all else equal, the relationship between the unknown variable and the long-run utility payoff is convex. A 50% chance of a 0% expropriation rate and a 50% chance of a 1% expropriation rate is preferable to a certain, compounding 0.5% expropriation rate.

Let us now consider the case of a philanthropist aiming to good by funding a more esoteric cause for which his endowment constitutes 50% of total current funding. We will hold the other parameters constant: the other funders discount at 1.5% per year, and so on.

Now, the patient philanthropist maximizes δ_P -discounted utility by investing until year $t^* = 46$ and then subsidizing the other funders so as to increase funding according to schedule

$$9.550 \cdot 10^8 \cdot e^{0.035(t-46)} - 9.475 \cdot 10^8 \cdot e^{0.025(t-46)} \quad (34)$$

for an expected δ_P -discounted utility payoff of

$$\int_{46}^{\infty} e^{-0.005t} \left(\ln \left(9.550 \cdot 10^8 \cdot e^{0.035(t-46)} \right) - \ln \left(9.475 \cdot 10^8 \cdot e^{0.025(t-46)} \right) \right) dt = 319. \quad (35)$$

An immediate transfer to the other cause's other funders would produce an expected δ_P -discounted payoff of

$$\int_0^{\infty} e^{-0.005t} (\ln(40B \cdot 0.015e^{0.025t}) - \ln(20B \cdot 0.015e^{0.025t})) dt = 138, \quad (36)$$

for a payoff ratio of $319/138 = 2.31$.

As we can see, the benefits of being strategic diminish sharply when the patient philanthropist is funding a cause for which he is one of the primary funders. Even in such cases, however, the optimal policy generally involves a long waiting period in which the patient philanthropist does not spend at all.

4 Model with one actor and varying effectiveness

As we have been assuming, a philanthropist whose spending on a cause constitutes a sufficiently large fraction of total spending on that cause will face diminishing returns to his philanthropic spending at any given time. He will thus want to smoothe his expenditures over time. He will also want to spend his funds when they will be most impactful, and he may expect impactfulness to vary arbitrarily over time.

In §2, we informally explored the donation-timing problem a patient philanthropist faces when he is a *marginal* donor to the cause of increasing human consumption, even after multiple centuries of compound interest. In §3, we explored the problem that arises when the patient philanthropist is one of a handful of uncoordinated funders of a cause, where the philanthropists have different discount rates, and the patient philanthropist is willing to invest until he becomes a substantial player. We will now explore the donation-timing problem a philanthropist faces when he is the *only* funder of a given cause. This assumption might be interpreted as an assumption about donor coordination and an assumption about the esotericism of the philanthropist’s utility function. The paradigmatic example will be the community of utilitarian-leaning philanthropists who believe that the most cost-effective interventions from a non-discounted utilitarian perspective are almost entirely neglected by society at large, and that these opportunities—opportunities to fund changes to the trajectory of human civilization with some probability, such as reductions in extinction risk, perhaps—are likely to be even more cost-effective in expected value than the trajectory-change of permanently increasing the growth path of future human consumption by investing or subsidizing investment. Conceiving of this community as a single philanthropist, we will explore how it should allocate its resources across time.

4.1 Hingeyness

Parfit (2011) famously hypothesized that he was living at the “hinge of history”. The hypothesis is that because we now have the opportunity to solve (or fail to solve) the risks of human extinction that have arisen over the past century, or because we can have a hand in setting up the technologies that will determine the course of future, the resources in our hands can be put to uses of much higher expected value than resources in the hands of our ancestors or descendants.

The image of a hinge evokes a single pivotal moment on which absolutely everything rests. Relaxing this claim somewhat, and translating it into the language of discounting, Parfit’s claim is the claim that a patient philanthropist’s discount rate in resources was negative in the past and is

positive in the future. A plot of a time series of “Resources’ leverage over the future” would look something like a bell curve, centered at 2011.

Without taking a stand as to the shape of the curve described above, we can certainly say that technological and historical developments make the metal of history more or less malleable. For lack of a better word for this malleability, let us somewhat playfully call it *hingeyness*.

Slightly more formally, we will say that a circumstance is “hingey” in proportion to the expected difference in value between the highest-expected-value use of resources and the expected value of continuing to invest them. Note that uncertainty is thus built into the definition: if buying a loaf of bread 500 years ago happened to cause a substantial shift to the course of civilization, but there was no way for the bread-buyer to know this, we will not say that her circumstances were hingey. Note also that hingeyness only tracks the *best* opportunities. Moments when it is cheap to do lasting damage, but difficult to prevent it, are unfortunate but not worth saving for.

Parfit did not mean that the hingeyest year was literally 2011, of course. But it is important to be precise about whether the times are currently getting more important or less. If Parfit was right down to the year, the patient philanthropist should currently spend his resources on high-impact opportunities quickly, since such opportunities are running out. On the other hand, if the most important year in this sense is 2061, the patient philanthropist should continue to invest—and perhaps get a few PhDs while he is at it—so that he can wield maximum influence when it is most important to do so.

Furthermore, note that a time-series of hingeyness would in fact exhibit peaks and valleys, in addition (if the “hinge of history hypothesis” is true) to a long-term upward trend. One important timing consideration concerns how big the peaks and valleys are in comparison to the long-term trend. The year 1941, for instance, was undoubtedly hingey. But consider a Brit living in 1941, with a few shillings to spare. Should she have put them toward the war effort, or do we think she would even then have done better in expected value to invest them so that they could help to safeguard and steer the development of artificial intelligence eighty years later? Or if a patient philanthropist currently in her twenties believes that the hingeyest year will be 2061, but also believes that her time is more valuable in 2019 than in the next few years because of temporary talent constraints in the pipeline of some effort of long-term importance, should she work or attend graduate school?

4.2 Basic model

Imagine that there is some measurable, empirical variable constructed so as to track the expected value, all things considered, of the future. It is what Bostrom (2014) would call an “evaluation heuristic”, or what I like to call a “longtermist economic index” (Greaves et al., 2019). It might consist of the power- or wealth-weighted distribution of moral values, the quality of the world’s technologies, the quality of our institutions, the quality of the natural environment, and the extent of peace. We will denote it by v .

We will say without loss of generality that the philanthropist has one unit of resources at time

$t = 0$. At each moment t , we will assume that the rate ρ at which the philanthropist can increase v is an isoelastic function, parametrized by η_P , of the rate x at which the philanthropist spends. (This η_P is unrelated to the consumption-welfare elasticity “ η ” in §2.) That is,

$$\rho_t(x_t) = h_t \frac{x(t)^{1-\eta_P}}{1-\eta_P} \quad (37)$$

where h_t (from “hingeyness”) is a scale parameter denoting how easy it is at time t to increase v .¹²

The philanthropist faces a constant instantaneous real interest rate r , a constant instantaneous discount rate δ_P (representing trends such as extinction risk, risks of “value drift”, any low rate pure time preference he may have, and other trends that might be roughly exponential), and a hingeyness schedule $h(t)$ (representing everything else). (We need not assume that r or δ_P is positive.) The fund’s problem is then to choose the continuous schedule of spending rates $x(t)$ that maximizes

$$\int_0^\infty e^{-\delta_P t} \rho(x(t)) dt = \int_0^\infty e^{-\delta_P t} h(t) \frac{x(t)^{1-\eta_P}}{1-\eta_P} dt \quad (38)$$

subject to the constraint

$$\int_0^\infty e^{-r t} x(t) dt = 1. \quad (39)$$

In the special case that $h(t)$ is constant, the constrained optimization is satisfied by

$$x(t) = \left(r - \frac{r - \delta_P}{\eta_P}\right) e^{\frac{r - \delta_P}{\eta_P} t}. \quad (40)$$

We have seen this formula before, in Proposition 3.1: it is just an infinite-horizon consumption-smoothing model under certainty, assuming either (a) no future outside income or (b) complete capital markets. (Recall that the assumption of complete capital markets renders this problem the same as the problem the philanthropist faces with no outside income. Given certainty and complete markets, the philanthropist with future income can borrow against his entire income stream.) (CITE relevant materials—maybe Friedman on the Permanent Income Hypothesis.) Nevertheless, given the centrality of the underlying relationship described above to the analysis below, let us now take a moment to note five of its features.

First: As one might expect, whether the outflows are increasing, constant, or decreasing in time depends on whether $r - \delta_P$ is greater than, equal to, or less than zero. Furthermore, observe the exponent on

$$x'(t) = \left(r - \frac{r - \delta_P}{\eta_P}\right) \left(\frac{r - \delta_P}{\eta_P}\right) e^{\frac{r - \delta_P}{\eta_P} t}. \quad (41)$$

If $r - \delta_P > 0$, the *rate of increase* in spending is also increasing with time, and if $r - \delta_P < 0$, $\lim_{x \rightarrow \infty} x(t) = 0$. It follows that we have no edge cases in which the fund grows or shrinks asymptotically to a positive size. The $r = \delta_P$ steady state is unstable.

¹²Note that since the size of the philanthropic fund is normalized to 1, the rate of money-spending x is given in “fraction of fund’s starting size per unit time”. A fund starting out at size m spending its money at rate x would therefore be able to push up v at the above rate multiplied by $m^{1-\eta_P}$; in other words, it would face a rescaled hingeyness schedule $\tilde{h} = m^{1-\eta_P} h$. So the above equation can represent the instantaneous utility function of an arbitrarily-sized fund, without loss of generality.

Second: As long as h is constant, the optimal spending schedule does not depend on h . This is due to the fact that h is, by construction, multiplicatively separable from the rest of our formula for $\rho(x)$. Per Footnote 11 above, it follows that if the fund faces isoelastic impact, the schedule on which it will spend money, expressed in fractions of its original size per unit time, will not depend on its original size.

Third: If h is a martingale, the optimal spending schedule is the same as if h is constant. Whatever value h may take at some time t , its expected value at all subsequent periods is equal to its observed value at t . It follows that the expected value of marginal spending at any given time is unchanged relative to the expected value of marginal spending at any other time.

Fourth: If $\eta_P = 1$, we have

$$x(t) = \delta_P e^{(r-\delta_P)t}. \quad (42)$$

The size of the fund at time t is thus

$$e^{rt} \left(1 - \int_0^t e^{-rs} \delta e^{(r-\delta_P)s} ds \right) = e^{rt} (2 - e^{-\delta_P t}). \quad (43)$$

We can now see that the rate we spend at time t as a proportion of fund size at time t equals

$$\frac{\delta_P e^{(r-\delta_P)t}}{e^{rt} (2 - e^{-\delta_P t})} = \frac{\delta_P e^{-\delta_P t}}{2 - e^{-\delta_P t}}, \quad (44)$$

which notably does not depend on r . More generally: if the interest rate goes up, the philanthropist faces both a price effect (as r goes up, future-spending becomes cheaper, so he wants to buy more of it) and an income effect (he will be spending so much more in the future anyway that, because of diminishing marginal impact, he finds less reason to invest so as to increase future spending). When $\eta_P < 1$, the more intuitive price effect predominates, and the investment rate increases in the interest rate. When $\eta_P > 1$, the income effect predominates, and investment decreases in the interest rate.

Fifth: If $\eta_P \geq 1$, a literal reading of the model implies that the impact of philanthropic spending is at least sometimes negative. This is illusory, since the model still implies that impact is larger (i.e. less negative) as the rate of spending increases. In other words, if the “impact of spending at rate x ” is negative, the “impact of spending at rate $x' < x$ ” is “even more negative”. If we would like to describe impact as a positive quantity, we can always do so by adding a sufficiently large positive constant to the impact function $\rho_t(x_t)$. This will not change any of our conclusions.

4.3 Markov-process model

We assumed in the basic model that hingeiness is constant over time. It is however precisely the claim that hingeiness varies over time, and that our time is uniquely hingeey, which leads so many to the conclusion that we ought to spend now rather than invest.

4.3.1 With arbitrary η_P

To capture (almost) any beliefs one might have about the ways in which hingeiness varies over time, therefore, let us allow h_t to follow an arbitrary finite Markov process. In doing so, we can also let

the model capture arbitrary finite beliefs about how δ_P and r might be expected to vary over time, and how these three variables might all be associated with each other.

The mechanism for doing this is simple. We have a finite set S of states s , each of which comes with a hingeiness value $h(s)$ and an interest rate $r(s)$. We also have a transition matrix T . We can leave out the state-specific discount rate $\delta_P(s)$, without loss of generality, by giving each state s transition probability $T_{s,s_E} = \delta_P(s)$ to an “extinction” attractor state s_E such that $h(s_E) = 0$ and $T_{s_E,s_E} = 1$.

Philanthropic resources will be divided optimally between spending and investing, in a given state, if the marginal value of spending equals the expected marginal value of investing.¹³ The expected marginal value of investing from a given state, in turn, is equal to the expected marginal value of money across subsequent states. And because money in a given state will be split between spending and investing, each of which will have equal marginal value, the expected value of money in each state is equal to the marginal value of additional spending in that state. We can therefore say that the philanthropist spends optimally in each state if *the marginal value of spending equals the expected marginal value of spending in the next period*.

Finally, remember that the isoelasticity assumption conveniently guarantees that the optimal proportion of the fund to spend in state s (or, in continuous time, the optimal rate at which to spend while in state s)—which we will denote $x(s)$ —depends only on η_P and on the features of s . In particular, it does not on the absolute size of the fund. To see this, observe that if the fund increases by proportion p , and the state-contingent spending policy stays the same, the marginal value of spending in each state falls by proportion $\eta_P p$. The marginal value of spending in each state will thus still equal the expected marginal value of spending in the next period; both quantities will fall by the same proportion.

More formally: for all $i \in \{1, \dots, |S|\}$,

$$h(s_i)(x(s_i))^{-\eta_P} = e^{r(s_i) - \delta_P(s_i)} \sum_{j=1}^{|S|} [T_{i,j} h(s_j) (e^{r(s_i)} (1 - x(s_i)) x(s_j))^{-\eta_P}] \quad (45)$$

The optimal spending function $x(s)$ is now given by $|S|$ equations with $|S|$ unknown variables (the $x(s_i)$), and it will generally be well-defined.

For ease of exposition, the above model description is set in discrete time. To find the optimal policy in continuous time, observe that the philanthropist spends at an *optimal rate* in each state if *the marginal value of spending in that state equals the expected marginal value of investing resources for the subsequent state*, when it arrives. We can find the latter by defining the transition matrix T such that $T_{i,j}$ is the instantaneous probability of a transition to j from i , with $T_{i,i} = 0 \forall i$ and $\sum_j T_{i,j}$, not necessarily summing to 1 for any i . Now, at any time in state i , the instantaneous probability

¹³The isoelastic functional form assumes infinite marginal utility when the spending rate is zero, so the philanthropist will never find himself in a corner solution of spending nothing or spending his entire budget. But this is of course a technicality.

of a transition to any other state is $\sigma(s_i) \triangleq \sum_j T_{i,j}$, and the probability density that the subsequent transition takes place t units of time into the future is given by $e^{-\sigma(s_i)t}$.

To find the expected value of investing marginal resources for the next state, whenever it arrives, we can therefore integrate over the possible transition times. Setting this expected value equal to the value of increasing the spending rate while in i , we have

$$h(s_i)(x(s_i))^{-\eta_P} = \int_0^\infty e^{-\sigma(s_i)t} \left(e^{(r(s_i) - \delta_P(s_i))t} \sum_j \left[T_{i,j} h(s_j) \left(e^{(r(s_i) - x(s_i))t} x(s_j) \right)^{-\eta_P} \right] \right) dt; \quad (46)$$

$$h(s_i)(x(s_i))^{-\eta_P} = \frac{\sum_j \left[T_{i,j} h(s_j) (x(s_j))^{-\eta_P} \right]}{\sigma(s_i) - r + \delta_P(s_i) + r\eta - x(s_i)\eta_P}. \quad (47)$$

4.3.2 With $\eta_P = 1$

The above setup lets us find the optimal spending policy numerically, and perhaps sheds some light on the shape of the problem. If $\eta_P = 1$ —that is, if impact is logarithmic in spending¹⁴—there is a simpler and more transparent way to see how the optimal spending schedule depends on the variables involved. In this case, if $h(s_i) \neq 0 \forall i$ (but observing that we can retain $\delta_P(s_i)$ directly), we can rearrange (47) so that we have $|S|$ linear equations of the form

$$\frac{1}{x(s_i)} = \frac{1}{\sigma(s_i) + \delta_P(s_i)} \left(\sum_j \left[T_{i,j} \frac{h(s_j)}{h(s_i)} \frac{1}{x(s_j)} \right] + 1 \right), \quad (48)$$

or a single matrix equation of the form

$$\vec{x} = \left(I_{|S|} - H \circ \mathcal{T} \right)^{-1} \vec{\sigma}, \quad (49)$$

where \vec{x} is the $|S|$ -vector of inverse spending proportions with $x_i = \frac{1}{x(s_i)}$, $\vec{\sigma}$ is the $|S|$ -vector of inverse instantaneous transition probabilities with $\sigma_i = \frac{1}{\sigma(s_i) + \delta_P(s_i)}$, H is the $|S| \times |S|$ “relative hingeyness matrix” with $H_{i,j} = \frac{h(s_j)}{h(s_i)}$, and \mathcal{T} is the $|S| \times |S|$ “relative transition probability matrix” with $\mathcal{T}_{i,j} = \frac{T_{i,j}}{\sigma(s_i) + \delta_P(s_i)}$.

4.3.3 Some numbers

Here is a simple application of the formula above. It is designed to test the force of the argument laid out in §2–3—namely that a patient philanthropist should typically invest—against the force of the argument that, if he finds himself at a particularly important moment, he ought to spend.

Suppose $\eta_P = 1$, $|S| = 2$, $\delta_P(s_1) = \delta_P(s_2) = 0.01$, $h(s_1) = 1$, $h(s_2) = 10$, $T_{1,2} = 0.1$, and $T_{2,1} = 0.8$. (Remember, r does not matter in the case of logarithmic impact.) In other words, suppose that

- the discount rate is relatively high, for someone with little or no pure time preference, at 1% per year;

¹⁴Find / cite piece from Owen Cotton-Barratt arguing that effort on problems of “unknown difficulty” will tend to offer roughly logarithmic returns.

- the ratio of hingeyness across states is high, at 10:1;
- hingey years are rare, with only a 10% chance of entering one from a non-hingey year; and
- hingey years are fleeting, with an 80% chance of leaving one given that one is in one.

Solving for \vec{x} , we find that $x(s_1) = 0.005$ and $x(s_2) = 0.048$. That is, even under the above parameters and in the very hingey state (s_2), the world’s patient philanthropists should still spend only 4.8% of their collective budget per year.

Under even more extreme parameter values, of course, the hingeyness effect can swamp the patience effect. But as this illustrates, the values have to be quite extreme before this happens. In sum, the case for waiting rather than giving appears to be stronger than intuition would have it, and stronger than longtermists in the EA community typically give it credit for.

5 Further considerations

5.1 Relationships to endogenize

The limitations of the timing model above are that η_P and T (and, implicitly, g) are exogenous and fixed and that the state space is finite. Exogeneity might be a reasonable assumption when the global pot of patient philanthropic funds is small enough—when the patient philanthropist can change the probability of who wins a world war, but not start a world war, say—but not when it is large. Generalizing the setup above to endogenize η_P , g , and T , and allowing for infinite states, would turn the patient philanthropist’s problem into the uselessly general problem of how to solve every possible Markov decision process. But perhaps there are some smaller, particularly important generalizations which would be useful and tractable.

5.1.1 Learning

Since learning enters the picture here only via h , our choice to make T exogenous also renders learning exogenous. Endogenous learning, and intervention reversibility, are conspicuously absent from the picture so far—especially given how centrally they feature in typical models of optimal timing, e.g. Pindyck (2002). This would of course be a very important extension, and this project is hopelessly incomplete without it.

5.1.2 The relationships between r , g , and fund size

[Informal notes:] Fund spending will grow more quickly than the growth rate g of the rest of the world economy iff $\frac{r-\delta_P}{\eta_P} > g$. But if the fund survives indefinitely long, it will eventually become a substantial fraction of the world economy. This will push down r and push up g .

- Note (ignoring h for now) that the fund will, if it acts optimally, just grow as a proportion of the world until $\frac{r-\delta_P}{\eta} = g$, not necessarily until it owns most of the world.

- As noted above, as time goes on, other patient actors will also come to own ever more of the world. (In fact this is already seems to be happening, at least to some extent (CITE).) (This will also push up g and push down r .) This means that even if the fund never spent anything, it would not grow to own everything; it would just grow to share the world with the other maximally-patient types. (This is relevant to the section below on “bargaining over the future.”)
- [For me to look into: evolutionary finance, the Cowley Rule.]

5.1.3 The relationship between δ_P and fund size

[Informal notes:] As the fund grows dizzyingly large, δ_P might also depend on its size (in absolute terms or, more likely, as a proportion of global assets). For instance, people might try to seize it—or it might grow better able to defend itself. Or its managers might be more inclined to succumb to value drift.

5.1.4 Uncertainty about long-run r , δ_P , and η_P

[Informal notes:]

- As noted above, the “Gollier part of the Gollier-Weitzman theorem”, as it’s sometimes called, is the common-sense observation that, over a long enough horizon, most of the value of being patient accrues in the scenarios where r is high and δ_P is low. So the more uncertain we are about the values of these rates, the more favorable it is to get the fund going (holding the expected rates constant).
- On the other hand, we should suppose that r will fall over the very long run; otherwise consumption would run into fundamental physical limits.
 - Longtermists sometimes model the growth process as falling from exponential to cubic once we (a) have reached technological maturity but (b) are engaging space colonization (see e.g. Tarsney, in progress).
 - We might expect δ_P to rise as we reach the heat death of the universe.
- At least over the very long run (and perhaps over the not-so-long run), we should expect η_P to change.

5.1.5 Bargaining over the future

[Informal notes:]

- So far we have assumed that the amount of influence the fund has on v at a given time t is simply an isoelastic function of how much it spends at that time, in absolute terms (multiplied by some black box, h_t). But what if the fund’s influence actually, or also, depends in some

way on what *fraction* of total spending it gets to direct? For instance, it could be that the course of the long-term future will be determined by who wins a big war between the patient philanthropist and a “patient misanthrope” in the year 3000, where the winner is whoever has more money to spend in the year 3000. In this case the value of saving another dollar is roughly equal to the impact of doing so on the probability that the philanthropist wins. There is no reason to think that this impact function is anything at all like isoelastic—or even concave.

- More realistic than a war, hopefully, is a big bargaining problem. To say anything detailed about how this plays out would seem to require speculating about the distribution of (a) utility functions among the “year 3000 patient-types” and (b) options available to them at the time, which it’s hard to see how one could possibly do. But we can still say some interesting / substantive things in this case. For example, if “fraction of the world owned” = “bargaining power”, and if the grand bargain plays out according to the axioms of Nash bargaining theory (i.e. if the outcome of a bargain between someone who owns fraction a of the world and someone who owns fraction $1 - a$ of the world is the same as the outcome of a Big Bargain among a crowd of identical people fraction a of whom have one goal and fraction $1 - a$ of whom have another, which seems reasonable), and if the goals among the funds are orthogonal and “rival”, then each fund’s payoff is linear in the fraction of the world it owns. We can then say things like:

- If the set of “potential things to do with the universe” is sparse (so that there’s basically no compromising to be done—two parties must settle on what one wants best, what the other wants best, or some lottery mixing the two), then, given an opportunity to grow not at its opponent’s expense, each fund’s utility function in resources will be hyperbolic. So in the limit, as the resource size increases, utility in additional(/lost) resources will be isoelastic with $\eta_P = 2$. [EXPLAIN.]
- It could be interesting to see what happens if we build in a rate c of moral convergence to (or divergence from) the correct moral theory over time, like the converse of value drift. The higher c is, the less useful it will be to own a larger fraction of the world in the future. On the other hand, the lower it is, the more likely it is that moral realism is false. So maybe a patient philanthropist under metaethical uncertainty should invest on the supposition that all the value is in the scenarios where c is sufficiently high, *even though* marginal future resources have less leverage in those scenarios.

5.2 Optimal investment

5.2.1 Mission hedging

[Informal notes:]

- Our construction of h , and our answer to the question of whether absolute or relative spending

is what matters, will inform how and to what extent the fund wants to “mission hedge” (Tran, 2017). I have not given this any deep thought. I do have the shallow and cynical thought that perhaps wars are the hingeiest times, such that the patient philanthropist should skew his investment portfolio toward weaponry...

- To the extent that this “bargaining” framing is right, where what really matters in the end is some concave function of what fraction of the world the patient philanthropist owns (rather than some concave function of the absolute amount), his ideal portfolio beta for any alpha is 1. Given the (presumably at least somewhat) positive relationship between beta and alpha, his optimal portfolio (other mission hedging concerns aside) will have beta a bit higher than 1. But it will not be quite so high as it would be in the naive case where he just cares about maximizing expected absolute return.

5.2.2 Movement building

[Informal notes:] It was noted in §2 that philanthropic expenditures with compounding flow-through effects are still best thought of as expenditures, rather than investments, because the resources they generate do not return to the philanthropist’s possession. An exception to this rule is “movement-building”: efforts aimed at growing the community of people who contribute to the philanthropist’s fund, or to funds with an identical objective.

How much a philanthropist should spend on movement-building is a complex problem. But if the philanthropist thinks (a) that resources put to purposes other than his own do no good, to a first approximation, and (b) that the good he can do is a function of his fund’s absolute size rather than of its size as a proportion of total wealth, then the problem simplifies. To determine how much to spend on movement-building efforts, the philanthropist then only needs to estimate the returns to various movement-building efforts and their covariance with the returns of other investments and movement-building efforts. The standard tools of modern portfolio theory should then allow him to construct the optimal portfolio of movement-building efforts and other investments.

5.2.3 Accounting for exogenous movement growth

[Informal notes:] Another reason sometimes given for spending now is that we might expect funding on the most important causes to increase exogenously over time. Cotton-Barratt (2015), for example, argues that those interested in mitigating existential risk should spend their resources directly on doing so now, even if they think that risk will increase over the coming century, because there are currently very few resources going to the cause but the flow appears to be increasing.

The simplest way to account for this possibility, I believe, is to model “the world at large” as a risky asset which every patient philanthropist already possesses.¹⁵ Money in others’ pockets is currently a near-worthless “asset” (in that almost none of it is being contributed to the world’s most important causes), but there is some chance that, in the future, it will become valuable (in that

¹⁵I thank Carl Shulman for this insight.

some non-negligible part of it may come to be spent on the world’s most important causes). In computing the patient philanthropist’s optimal investment portfolio, we should just throw this asset into the mix.

6 Conclusion

It is often argued that, as patient philanthropists in an impatient world, we should expect *a priori* that the most neglected (and therefore highest-expected-value) causes and interventions are those most of whose value accrues in the distant future. From here it is most commonly argued that we should expect that the most underfunded efforts are those aimed at reducing “existential risk”—that is, risks of human extinction, or other risks with similarly long-term negative consequences (see e.g. Bostrom 2003). Some (e.g. Beckstead 2013) have argued that, more generally, what we should expect to be most neglected are trajectory changes most of whose total value accrues in the distant future, of which reductions in existential risk are just one possible kind.

By this line of reasoning, there are interventions we should expect *a priori* to be even more underfunded, in an impatient world, than efforts to reduce existential risk or effect similarly long-lasting trajectory changes. These are efforts to reduce existential risk and effect trajectory change which *do not begin for a long time*. However much the world should invest to spend on lowering the existential risks of the next millennium, we should expect that the world is underfunding this potential investment even more severely than it underfunds the risk mitigation efforts of today—even after accounting for the uncertainty that comes with planning for the future, and the fact that we might expect people in the future to spend substantially on reducing the existential risks that confront them.

Perhaps we face temptingly cost-effective opportunities to produce persistent improvements to others’ welfare. Perhaps we find ourselves at a particularly hinge time. Neither of these considerations settle the question of whether to give or to invest. The future, too, will have its opportunities for charity. As patient philanthropists, we should not neglect them.

7 References

- [1] Alexander, S. (2013). “Investment and Inefficient Charity”, SlateStarCodex.
URL: <https://slatestarcodex.com/2013/04/05/investment-and-inefficient-charity>
(accessed 1 Jun 2019).
- [2] Barrage, L. (2018). “Be careful what you calibrate for: Social discounting in general equilibrium”, *Journal of Public Economics* 160: 33–49.
- [3] Beckstead, N. (2013). “A Proposed Adjustment to the Astronomical Waste Argument”, LessWrong. URL: http://lesswrong.com/lw/hjb/a_proposed_adjustment_to_the_astronomical_waste (accessed 30 May 2019).

- [4] Bostrom, N. (2003). “Astronomical Waste: The Opportunity Cost of Delayed Technological Development”, *Utilitas* 15(3): 308–14.
- [5] Bostrom, N. (2014). “Crucial Considerations and Wise Philanthropy”. URL: <https://www.effectivealtruism.org/articles/crucial-considerations-and-wise-philanthropy-nick-bostrom/> (accessed 1 Jun 2019).
- [6] Bradford, D. F. (1975). “Constraints on Government Investment Opportunities and the Choice of Discount Rate”, *The American Economic Review* 65(5): 887–99.
- [7] Broome, J. (1994). “Discounting the Future”, *Philosophy and Public Affairs* 23(2): 128–56.
- [8] Christiano, P. F. (2013). “Giving now vs. later”, Rational Altruist. URL: <https://rationalaltruist.com/2013/03/12/giving-now-vs-later> (accessed 30 May 2019).
- [9] Christiano, P. F. (2014). “We can probably influence the far future”, Rational Altruist. URL: <https://rationalaltruist.com/2014/05/04/we-can-probably-influence-the-far-future> (accessed 30 May 2019).
- [10] Cotton-Barratt, O. (2015). “Allocating risk mitigation across time”, Future of Humanity Institute Technical Report #2015-2.
- [11] Credit Suisse Research Institute (2018). “Global Wealth Report 2018”. URL: <https://www.credit-suisse.com/media/assets/corporate/docs/about-us/research/publications/global-wealth-report-2018-en.pdf> (accessed 29 May 2019).
- [12] Dickens, M. (1994). “Should Global Poverty Donors Give Now or Later? An In-Depth Analysis”, Philosophical Multicore. URL: https://mdickens.me/2019/01/21/should_global_poverty_donors_give_now_or_later (accessed 1 Jun 2019).
- [13] Drupp, M. A., M. C. Freeman, B. Groom, and F. Nesje (2018). “Discounting Disentangled”, *American Economic Journal: Economic Policy* 10(4): 109–34.
- [14] Giglio, S., M. Maggiori, and J. Stroebel (2015). “Very Long-Run Discount Rates”, *The Quarterly Journal of Economics* 130(1): 1–53.
- [15] Greaves, H., W. MacAskill, R. O’Keeffe-O’Donovan, and P. Trammell (2019). “A Research Agenda for the Global Priorities Institute”. URL: <https://globalprioritiesinstitute.org/wp-content/uploads/gpi-research-agenda.pdf> (accessed 1 Jun 2019).
- [16] Giving What We Can (2013). “Frequently Asked Questions”. URL: <https://www.givingwhatwecan.org/about-us/frequently-asked-questions> (accessed 29 May 2019).
- [17] Gollier, C. (2016). “Evaluation of long-dated assets: The role of parameter uncertainty”, *Journal of Monetary Economics* 84: 66–83.

- [18] Gollier, C. and M. Weitzman (2010). “How should the distant future be discounted when discount rates are uncertain?”, *Economics Letters* 107(3): 350–3.
- [19] Hanson, R. (2013). “If More Now, Less Later”, *Overcoming Bias*.
URL: <http://www.overcomingbias.com/2013/04/more-now-means-less-later.html>
(accessed 28 May 2019).
- [20] Harris, K. (2018). Unpublished.
- [21] Kaldor, N. (1957). “A Model of Economic Growth”, *The Economic Journal* 67(268): 591–624.
- [22] Karnofsky, H. (2007). “More thoughts on ‘responsible investing’”, *The GiveWell Blog*.
URL: <https://blog.givewell.org/2007/01/22/more-thoughts-on-responsible-investing>
(accessed 29 May 2019).
- [23] MacAskill, W. (2015). *Doing Good Better* (New York: Random House).
- [24] Mankiw, N. G. (2015). “Yes, $r > g$. So what?”, *American Economic Review: Papers & Proceedings* 105(5): 43–7.
- [25] Olson, M., and M. Bailey (1981). “Positive Time Preference”, *Journal of Political Economy* 89(1): 1–25.
- [26] Ord, T. (2017). “The value of money going to different groups”. The Centre for Effective Altruism blog. URL: <https://www.centreforeffectivealtruism.org/blog/the-value-of-money-going-to-different-groups> (accessed 31 May 2019).
- [27] Parfit, D. (2011). *On What Matters, Vol. 2* (Oxford: Oxford University Press).
- [28] Piketty, T. (2014). *Capital in the Twenty-First Century*, translated by A. Goldhammer (Cambridge, MA: Harvard University Press).
- [29] Pindyck, R. S. (2002). “Optimal Timing Problems in Environmental Economics”, *Journal of Economic Dynamics & Control* 26: 1677–97.
- [30] Ramsey, F. (1928). “A Mathematical Theory of Saving”, *The Economic Journal* 38(152): 543–59.
- [31] Tarsney, C. (2019). “The Epistemological Challenge to Longtermism”. Unpublished.
- [32] Tran, B. R. (2016). “Divest, Disregard, or Double Down?”, *Federal Reserve Board Finance and Economics Discussion Series* 2017-42.
- [33] Vivalt, E. (2019). “Give Later”.
URL: <http://evavivalt.com/2019/07/give-later>
(accessed 13 Aug 2019).