

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 1.1 | Academic overview | 1 |
| 1.2 | Philanthropic overview | 3 |
| 2 | Basic model | 4 |
| 2.1 | Model | 4 |
| 2.2 | Discussion | 5 |
| 3 | Model with varying parameters | 6 |
| 3.1 | General Markov-process model | 6 |
| 3.2 | Simplification with $\eta = 1$ | 8 |
| 3.3 | Expected payoffs | 9 |
| 4 | Interactions with impatient actors | 10 |
| 4.1 | Motivation and framework | 10 |
| 4.2 | Spending without commitment | 11 |
| 4.3 | Spending with commitment | 11 |
| 4.4 | Linear subsidization | 13 |
| 4.5 | Examples | 21 |
| 5 | Applications to philanthropic cause-areas | 23 |
| 5.1 | Direct efforts to increase near-term human welfare | 23 |
| 5.1.1 | Public projects and the Ramsey formula: “Discount by $\delta + \eta g$ ” | 23 |
| 5.1.2 | A common mistake: “Discount by $\eta g^* > \delta + \eta g$ ” | 25 |
| 5.1.3 | A proposed correction: “Discount by $\eta g < \delta + \eta g$ ” | 26 |
| 5.2 | Longtermist efforts | 28 |
| 6 | Further considerations | 31 |
| 6.1 | Relationships to endogenize | 31 |
| 6.1.1 | Endogenous learning | 31 |
| 6.1.2 | The relationships between r , g , and fund size | 31 |
| 6.1.3 | The relationship between δ_P and fund size | 32 |
| 6.1.4 | Uncertainty about long-run r , δ_P , and η | 33 |
| 6.1.5 | Bargaining over the future | 33 |
| 6.2 | Investment strategies | 35 |
| 6.2.1 | Mission hedging | 35 |
| 6.2.2 | Movement building | 35 |
| 6.2.3 | Accounting for exogenous movement growth | 35 |
| 6.2.4 | Leverage | 36 |
| 7 | Conclusion | 36 |
| 8 | References | 38 |

Discounting for Patient Philanthropists

Philip Trammell*

July 19, 2020

EARLY DRAFT

Abstract

Philanthropists must decide to what extent to spend their resources on present philanthropic projects and to what extent to invest them for use on future philanthropic projects. Furthermore, when a philanthropist aims to provide public goods for which he is not the only funder—and, in particular, when he and the other funders have different rates of pure time preference—he must consider the ways in which his spending schedule affects that of the good’s other funders. I investigate some features of the resulting discounting problem in a relatively general setting in which the cost of the public good, the interest rate, and the discount rate can vary and covary arbitrarily over time, and in which the patient philanthropist has the opportunity to subsidize future spending by his less patient partners (if any). In the presence of other funders, I find that the patient philanthropist often does best to exclusively invest he is wealthy enough to fund subsidies which implement the patient-optimal funding schedule, and that the relative gains to doing so increase without bound as the relative wealth of the patient philanthropist decreases to zero. Finally, I explore how the model interacts with discussions among philanthropists regarding the optimal timing of (a) direct efforts to improve near-term human welfare and (b) efforts aimed more directly at increasing the expected value of the long term. In both cases, I conclude that standard assumptions imply that patient philanthropists should invest most of their resources in most circumstances.

1 Introduction

1.1 Academic overview

We will call someone a “philanthropist” if he aims to contribute to the provision of a public good. We will call someone “patient” if he has low (including zero) pure time preference with respect to the

*Global Priorities Institute and Department of Economics, University of Oxford. Contact: philip.trammell@economics.ox.ac.uk. Thanks to Lewis Ho, Hilary Greaves, Andreas Mogensen, Rossa O’Keeffe-O’Donovan, Frikk Nesje, Bastian Stern, Lennart Stern, Carl Shulman, Christian Tarsney, Benjamin Tereick, and Teruji Thomas for helpful corrections and comments. All remaining errors and omissions are my own.

welfare he creates by providing a good. 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. Many philanthropists, including many engaged with the “Effective Altruism” (EA) community, endorse this position. Nevertheless, it is currently standard practice for the providers of most public goods—including governments (see e.g. Council of Economic Advisers, 2017), many philanthropists, and individuals (see e.g. Olson and Bailey, 1981) whose private spending carries some positive externality—to act roughly so as to maximize the welfare they provide their beneficiaries after discounting by some positive and substantial rate of pure time preference. In determining the schedule on which to disburse his assets, a patient philanthropist must therefore consider his impact not only on the direct provision of the good in question but also on the behavior of the good’s other, less patient providers.

In §2, I observe that the public good provision timing problem faced by a single funder, regardless of whether this funder is a policymaker or a philanthropist with an unusually low discount rate, is analogous to a household’s basic consumption-smoothing problem. The rest of the paper builds on this consumption-smoothing framework in three ways.

First, while existing work has largely focused on the patient-optimal provision of public goods in the relatively narrow contexts of abating climate change or increasing future domestic consumption, I examine the problem in less detail but a more general setting. Disagreements among philanthropists regarding the optimal spending rate sometimes depend on idiosyncratic disagreements about future interest rates; philanthropic projects’ “hazard rates”; costs of doing good, in domains not limited to climate or human consumption; and/or how these variables associate with one another. As laid out in §3, therefore, I take as input an arbitrary finite state space, with state-specific interest rates, hazard rates, and good costs, and I allow movement in the space to follow an arbitrary transition matrix in continuous time. This framework allows the reader to determine how optimally to smooth philanthropic spending given her own philanthropic beliefs and preferences.

Second, in §4 I apply the model to settings in which a public good has multiple funders with differing discount rates, so that patient philanthropists must worry about the extent to which their saving crowds out others’ saving. In one respect this project resembles existing work on private contributions to public goods (e.g. by Bergstrom et al., 1986), in which contributors must worry about the extent to which their contributions crowd out others’. I address inter- rather than intra-temporal crowding out, however, and so assume that the net present value of the budget each party has allocated to the public good is fixed. In another respect the project resembles existing work on optimal taxation by policymakers with lower discount rates than their constituents, as done by Barrage (2018), von Below (2012), Belfiori (2017), Farhi and Werning (2005, 2010), and others. The premise of this literature is that patient policymakers might like to invest for future spending, but that to avoid crowding out private investment, they must instead subsidize private investment and tax private consumption. Since I am addressing a patient philanthropist rather than a patient policymaker, however, I focus on the constraint that only subsidization, and not taxation, is feasible.

Third, in §5 I discuss the relationship between the conclusions from the earlier sections and the

existing discussion about donation timing produced by philanthropists outside academia, primarily those within the EA movement.

I conclude that the case for currently investing most or all “patient philanthropic resources” is stronger than it would appear from existing discussion, regardless of whether such funds are ultimately best spent on interventions to increase human welfare in the near term or on efforts most of whose value is intended to accrue in the distant future. I argue that this conclusion would hold even in a single-funder scenario, and that the presence of impatient funders further strengthens the case that patient philanthropists should invest.

1.2 Philanthropic overview

As noted above, a central conclusion of the model presented here is that, according to standard economic assumptions and seemingly plausible assumptions about our ability to achieve long-lasting impact, patient philanthropists should invest, rather than spend, most or all of their resources in most circumstances. This conclusion would appear to recommend a substantial departure from current common practice.

The EA movement, for example, brands itself in part as a community of impartial—and by extension patient—philanthropists. Its members include Dustin Moskovitz and Cari Tuna, who plan to give away almost all of their roughly \$11.1B net worth (Matthews, 2018), and numerous other high-net-worth individuals. Members of Giving What We Can, an EA-affiliated nonprofit that encourages individuals to give at least ten percent of their incomes to effective charities, have also collectively pledged an estimated \$1.5 billion, mostly from smaller donors. There is currently no consensus within the movement regarding whether philanthropic funds are best spent, from an impartial perspective, on direct efforts to improve others’ wellbeing in the near term or on more speculative efforts to improve the distant future (Moss et al., 2019). Nevertheless, there appears to be a tendency across cause areas to conclude that one does most good by giving now—and implicitly, given how little attention the question has recently received, 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 (Giving What We Can, 2013). For efforts focused on the long term, it is often argued that we are currently living at a special time in history, during which our efforts to influence the future are uniquely impactful (see e.g. Parfit, 2011). This tendency is not unanimous: a few have argued that investment returns exceed the rate at which doing good is growing more costly (see e.g. Hanson, 2013), and some appear recently to be rethinking their old positions for giving now (see e.g. Vivalta, 2019). But advocating long-term investment is still far from mainstream in philanthropic circles, including the EA movement, and very little effort has been put into its implementation.

Perhaps ironically, 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). Most small donors outside the EA community, however, also appear to spend their charity budgets as soon as they earn them. This is difficult to measure precisely, of course, but note that individual giving by Americans totaled \$427.71B in 2018 (Giving USA, 2019), whereas contributions by Americans to donor-advised funds—tax-exempt vehicles in which donors without private foundations can invest funds for future giving—totaled \$37.12B in 2018 (National Philanthropic Trust, 2019).

In short, if it can be demonstrated that patient philanthropists should spend slowly, these conclusions could be highly relevant to the spending of at least tens of billions of dollars. A careful attempt to model the philanthropist’s discounting problem therefore appears to be worthwhile.

2 Basic model

2.1 Model

Let us begin with a highly simplified model in which an agent is the sole provider of some good, and let us denote the size of the agent’s budget at time $t = 0$ by B . At each moment t , we will assume that the flow utility u achieved by providing the good is an isoelastic function, parametrized by η , of the rate x at which the agent spends. That is,¹

$$u(x(t)) = \frac{x(t)^{1-\eta}}{1-\eta}. \quad (1)$$

The agent faces a constant instantaneous real interest rate r and a constant instantaneous discount rate δ . (For a patient philanthropist, this might represent factors such as the annual risk that his funds are expropriated, that he or his fund’s inheritors drift from his original values, or that the world ends; any low rate pure time preference he may have; and other trends that might roughly exponentially diminish the value of waiting to spend). We need not assume that r or δ is positive. The agent’s problem is then to choose the schedule of spending rates $x(t)$ that maximizes

$$\int_0^\infty e^{-\delta t} u(x(t)) dt = \int_0^\infty e^{-\delta t} \frac{x(t)^{1-\eta}}{1-\eta} dt \quad (2)$$

subject to the constraint

$$\int_0^\infty e^{-rt} x(t) dt \leq B. \quad (3)$$

Proposition 2.1 (*The optimal individual spending schedule under constant parameters*)

Suppose an individual (philanthropic or otherwise) has isoelastic utility in spending parameterized by η , a constant discount rate δ , and a budget B , and suppose she can invest her resources at a constant interest rate r . Then the individual maximizes discounted utility by following spending schedule

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

¹Isoelastic utility is defined as $u(x(t)) = \ln(x)$ when $\eta = 1$, by a limit condition.

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

$$\frac{\partial}{\partial[y(t)]} \left[e^{-\delta t} \frac{(e^{rt}y(t))^{1-\eta}}{1-\eta} \right] = k \quad \forall t, \quad \eta \neq 1; \quad (4)$$

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

Taking the derivative and rearranging, we have

$$y(t) = k^{\frac{-1}{\eta}} e^{\frac{r-r\eta-\delta}{\eta}t}. \quad (5)$$

Subjecting this resource allocation to the budget constraint, we have

$$\int_0^\infty k^{\frac{-1}{\eta}} e^{\frac{r-r\eta-\delta}{\eta}t} dt = B; \quad (6)$$

$$k = \left(B \left(\frac{r\eta - r + \delta}{\eta} \right) \right)^{-\eta}. \quad (7)$$

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

$$x(t) = B \left(\frac{r\eta - r + \delta}{\eta} \right) e^{\frac{r-\delta}{\eta}t} \quad (8)$$

as desired. ■

Corollary 2.2 (*The payoff under constant parameters*)

Following the optimal spending schedule $x(t)$, as given in Proposition 2.1, produces a payoff of

$$\frac{\eta \left(\frac{B}{\eta} \right)^{1-\eta}}{(1-\eta)(r\eta - r + \delta)\eta}, \quad \eta \neq 1;$$

$$\frac{\delta \ln(B\delta) + r - \delta}{\delta^2}, \quad \eta = 1.$$

This can be calculated straightforwardly from the integral $\int_0^\infty e^{-\delta t} u(x(t)) dt$. ■

2.2 Discussion

The above model is motivated in this paper by the scenario in which a philanthropist is the sole provider of some public good. So far, it is equivalent to an infinite-horizon consumption-smoothing model under certainty, assuming either (a) no future outside income or (b) complete capital markets. (Note that the assumption of complete capital markets renders this problem the same as the problem one faces with no outside income. Given certainty and complete markets, someone with future income can borrow against her entire income stream, and B can represent current assets plus the present value of future income.) Nevertheless, given the centrality of the underlying relationship described

above to the analysis of patient philanthropy below, let us now take a moment to note three of its relevant features.

First, and most importantly: In the context of a simple consumption-smoothing model, the optimal spending rate is highly sensitive to the discount rate. The patient, that is, should spend slowly. As we can see from (8) at $t = 0$, it is always optimal to spend at proportional rate $\frac{r\eta - r + \delta}{\eta}$. In particular, if $\eta = 1$, the spending rate should equal δ . For instance, philanthropists who are funding idiosyncratic projects with no other present or future funders, who discount future impacts at 0.1% per year, and who are confident that the world (or their philanthropic projects) will not soon be brought to an end, should spend only 0.1% of their budgets per year.

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

$$x'(t) = B \left(\frac{r\eta - r + \delta}{\eta} \right) \left(\frac{r - \delta}{\eta} \right) e^{\frac{r - \delta}{\eta} t}. \quad (9)$$

If $r - \delta > 0$, the *rate of increase* in spending is also increasing with time, and if $r - \delta < 0$, $\lim_{x \rightarrow \infty} x(t) = 0$. It follows that we have no edge cases in which one's assets should be expected to grow or shrink asymptotically to a positive size. The $r = \delta$ steady state is unstable. If a fund should grow, it should not stop growing, even once it has grown very large.

Third: If $r\eta - r + \delta \leq 0$, we appear to be led to the conclusion that it is always preferable to invest than to spend. This is, in other words, the circumstance which gives rise to Koopmans' (1965, 1967) "paradox of the indefinitely postponed splurge". For the purposes of this paper, rather than broach the subject of infinite ethics, we will assume that this circumstance does not hold. Note that it does not whenever $\eta > 1$, $r > 0$, and $\delta \geq 0$. That is, under the reasonable assumption that $\eta > 1$, agents can be arbitrarily patient without approaching paradoxical territory.

3 Model with varying parameters

3.1 General Markov-process model

The basic consumption-smoothing model introduced above, according to which patient philanthropists should spend slowly, assumes that the interest rate, discount rate, and flow utility achieved by a given expenditure rate are constant over time. As noted in the introduction, however, some patient philanthropists advocate for a higher spending rate on the grounds that they expect these parameters to vary. We will account for these complications as generally as possible by introducing a finite set S of states s , each of which comes with a state-specific interest rate $r(s)$, discount rate $\delta(s)$, and scale parameter $h(s)$, such that the flow utility from spending at rate x in state s is given by

$$u_s(x) = h(s) \frac{x^{1-\eta}}{1-\eta}. \quad (10)$$

Instantaneous transition probabilities among the states will be given by a transition matrix T .

Note that this is equivalent to a model in which the *price* of the good being provided is a state-dependent value $p(s)$, with $h(s) \triangleq p(s)^{\eta-1}$.

Relatedly, note that as long as h is constant, the optimal spending schedule does not depend on h . More generally, because h is by construction multiplicatively separable from the rest of our formula for $u(x)$, rescaling $h(s)$ for all s amounts to a change of units, and does not affect the relative value of spending across periods.

Finally, note that if h is a martingale, the optimal spending schedule is the same as if h is constant. In this case, 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.

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.² 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 resources 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 agent spends optimally in each state if the agent's intertemporal Euler equation is satisfied—that is, if the marginal value of spending equals the expected marginal value of investing to spend in the next period.

The isoelasticity assumption conveniently guarantees that the optimal proportion of funds 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 and on the features of s . In particular, it does not depend on the absolute size of the fund. To see this, observe that if the fund increases by some proportion m , and the state-contingent spending policy stays the same, the marginal value of spending in each state falls by proportion ηm . 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:

Proposition 3.1 (*The optimal spending policy given varying parameters in discrete time*)

Suppose an individual has isoelastic utility in spending parameterized by η and a state-dependent discount rate $\delta(s)$, and suppose she can invest her resources at a state-dependent interest rate $r(s)$. Suppose also that movement in the state space S is given by transition matrix T . Then the individual

²The isoelastic functional form assumes infinite marginal utility when the spending rate is zero, so the agent will never quite find himself in a corner solution of spending nothing or spending his entire budget.

maximizes discounted utility by following spending policy $x(s)$ such that, for all $i \in \{1, \dots, |S|\}$,

$$h(s_i)(x(s_i))^{-\eta} = e^{r(s_i) - \delta(s_i)} \sum_{j=1}^{|S|} [T_{i,j} h(s_j) (e^{r(s_i)} (1 - x(s_i)) x(s_j))^{-\eta}].$$

The optimal spending policy 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 is set in discrete time. To find the optimal policy in continuous time, observe that the agent 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 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 $\sigma(s_i) e^{-\sigma(s_i)t}$.

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

$$h(s_i)(x(s_i))^{-\eta} = \int_0^\infty \sigma(s_i) e^{-\sigma(s_i)t} \left(e^{(r(s_i) - \delta(s_i))t} \sum_j \left[\frac{T_{i,j}}{\sigma(s_i)} h(s_j) \left(e^{(r(s_i) - x(s_i))t} x(s_j) \right)^{-\eta} \right] \right) dt. \quad (11)$$

Simplifying this expression, we have the following:

Proposition 3.2 (*The optimal spending policy given varying parameters in continuous time*)

Suppose an individual has isoelastic utility in spending parameterized by η and a state-dependent discount rate $\delta(s)$, and suppose she can invest her resources at a state-dependent interest rate $r(s)$. Suppose also that movement in the state space S is given by Poisson processes with transition rates from s_i to s_j given by matrix T . Then the individual maximizes discounted utility by following spending policy $x(s)$ such that, for all $i \in \{1, \dots, |S|\}$,

$$h(s_i)(x(s_i))^{-\eta} = \frac{\sum_j [T_{i,j} h(s_j)(x(s_j))^{-\eta}]}{\sigma(s_i) - r(s_i) + \delta(s_i) + r(s_i)\eta - x(s_i)\eta}.$$

3.2 Simplification with $\eta = 1$

The above setup lets us find the optimal spending policy numerically, and perhaps sheds some light on the shape of the problem. If $\eta = 1$, however—that is, if impact is logarithmic in spending—there is an analytic solution, offering perhaps a simpler and more transparent look into how the optimal spending schedule depends on the variables involved.

Proposition 3.3 (*The optimal spending policy given varying parameters in continuous time, where $\eta = 1$*)

Suppose an individual has logarithmic utility in spending and a state-dependent discount rate $\delta(s)$, and suppose she can invest her resources at a state-dependent interest rate $r(s)$. Suppose also that movement in the state space S is given by Poisson processes with transition rates from s_i to s_j given by matrix T . Then the individual maximizes discounted utility by following the spending policy given by matrix equation

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

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 $\sigma_i = \frac{1}{\sigma(s_i) + \delta(s_i)}$, H is the $|S| \times |S|$ “scaling matrix” with $H_{i,j} = \frac{h(s_j)}{h(s_i)}$, and T is the $|S| \times |S|$ “relative transition probability matrix” with $T_{i,j} = \frac{T_{i,j}}{\sigma(s_i) + \delta(s_i)}$.

The result follows from substituting $\eta = 1$ into the expression in Proposition 3.2, so that we have $|S|$ linear equations of the form

$$\frac{1}{x(s_i)} = \frac{1}{\sigma(s_i) + \delta(s_i)} \left(\sum_j \left[T_{i,j} \frac{h(s_j)}{h(s_i)} \frac{1}{x(s_j)} \right] + 1 \right), \quad (12)$$

and rearranging. ■

3.3 Expected payoffs

Proposition 3.4 (*The expected payoff given varying parameters in continuous time*)

Let

$$a_i \triangleq 1/(\sigma(s_i) + \delta(s_i) + (x(s_i) - r(s_i))(1 - \eta)).$$

If $\eta \neq 1$, the expected payoff v_i by state s_i from following a spending policy $x(s)$ is given by

$$\vec{v} = B^{1-\eta} (I_{|S|} - A \circ T)^{-1} \vec{z},$$

where A is an $|S| \times |S|$ matrix with $A_{i,j} \triangleq a_i$ and \vec{z} is an $|S|$ -vector with

$$z_i \triangleq h(s_i) \frac{x(s_i)^{1-\eta}}{(1-\eta)} a_i.$$

If $\eta = 1$, the expected payoff v_i by state s_i from following a spending policy $x(s)$ is given by

$$\vec{v} = (I_{|S|} - A \circ T)^{-1} (\vec{z} + \vec{w}_\delta \ln B),$$

where a_i and A are as defined above, \vec{z} is an $|S|$ -vector with

$$z_i \triangleq \left(\frac{r(s_i) - x(s_i) + \delta(s_i) \ln x(s_i)}{\delta(s_i)} + \frac{\sigma(s_i) h(s_i) (x(s_i) - r(s_i))}{\delta(s_i) (\sigma(s_i) + \delta(s_i))} \right) a_i + \frac{\sigma(s_i)}{\sigma(s_i) - r(s_i) + x(s_i) + \delta(s_i)},$$

and \vec{w}_δ is an $|S|$ -vector with $w_{\delta,i} \triangleq \sigma(s_i) a_i / \delta(s_i)$.

The result follows straightforwardly from two observations. First, observe that, in each state s_i , the expected payoff must equal the utility that will accrue from spending at rate $x(s_i)$ until the state transition, plus the expected payoff following the transition. Second, observe that multiplying one's budget by some constant m multiplies spending at all times by m .

If $\eta \neq 1$, multiplying one's budget by m thereby multiplies flow utility at all times, and thus the payoff, by $m^{1-\eta}$. Given a starting budget of 1 and a transition time of t , for example, the continuation payoff at t will equal $e^{(r(s_i)-x(s_i))(1-\eta)t}$ multiplied by the continuation payoff in the new state given a budget at t of 1. Our payoffs must therefore satisfy the $|S|$ equations given by

$$v(s_i) = B^{1-\eta} \int_0^\infty \sigma(s_i) e^{-\sigma(s_i)t} \left[\int_0^t e^{-\delta(s_i)q} h(s_i) \frac{(x(s_i) e^{(r(s_i)-x(s_i))q})^{1-\eta}}{1-\eta} dq \right. \\ \left. + e^{-\delta(s_i)t} \sum_{j=1}^{|S|} \frac{T_{i,j}}{\sigma(s_i)} e^{(r(s_i)-x(s_i))(1-\eta)t} v(s_j) \right] dt. \quad (13)$$

If $\eta = 1$, multiplying one's budget by m thereby adds $\ln m$ to flow utility at all times. We can therefore find the payoff given budget B by finding the payoff v_1 given budget 1 and adding it to the expected discounted value v_B of a stream of payoffs of size $\ln B$. These two terms will respectively satisfy

$$v_1(s_i) = \int_0^\infty \sigma(s_i) e^{-\sigma(s_i)t} \left[\int_0^t e^{-\delta(s_i)q} h(s_i) \ln(x(s_i) e^{(r(s_i)-x(s_i))q}) dq \right. \\ \left. + e^{-\delta(s_i)t} \sum_{j=1}^{|S|} \frac{T_{i,j}}{\sigma(s_i)} (e^{(r(s_i)-x(s_i))t} + v(s_j)) \right] dt \quad (14)$$

and

$$v_B(s_i) = \int_0^\infty \sigma(s_i) e^{-\sigma(s_i)t} \left[\int_0^t e^{-\delta(s_i)q} \ln B dq + e^{-\delta(s_i)t} \sum_{j=1}^{|S|} \frac{T_{i,j}}{\sigma(s_i)} v_B(s_j) \right] dt \quad (15)$$

Simplifying and rearranging these expressions gives the results. ■

4 Interactions with impatient actors

4.1 Motivation and framework

The model above allows us to determine the optimal spending policy regarding the provision of a good for which there is only one purchaser. It applies, for instance, to the schedule on which an individual should allocate her private spending. When one is a philanthropist providing a public good, however, one must consider the ways in which one's own funding affects the behavior of the good's other funders. In particular, when one is a patient philanthropist, one must remember that investment for future spending will induce less patient funders to spend more quickly.

This section will not begin to constitute an exhaustive account of how the patient and the impatient might be expected to interact. As we will see, however, across a variety of scenarios, intertemporal crowd-out concerns can motivate substantially different—and, in general, even “more

patient”—behavior from a patient philanthropist than is optimal in the single-funder context. In particular, the patient philanthropist often does best in the presence of impatient funders to invest *all* his resources, for some period, and then to spend on an exponential schedule resembling the single-funder schedule determined above.

Throughout the results below, we will refer to the impatient funder(s) of the public good in question as “the beneficiary” and to the patient funder(s) as “the philanthropist”. We will conceive of this good as the beneficiary’s “consumption”. We will denote the beneficiary’s budget at time t by B_t , with $B \triangleq B_0$; the philanthropist’s budget at time t by $B_{P,t}$ with $B_P \triangleq B_{P,0}$; the impatient funder’s state-dependent discount rate by $\delta(s)$; and the philanthropist’s discount rate by $\delta_P(s) = \delta(s) - \rho$, where $\rho > 0$ might be interpreted as the portion of the beneficiary’s discount rate consisting of pure time preference.³ At every moment $t \geq 0$, the players observe the spending and state histories and independently⁴ choose spending rates x and x_P respectively.

4.2 Spending without commitment

In the presence of a strategic beneficiary, the philanthropist can do substantially better by acting strategically than by acting naively. To see this, we will find a worst subgame-perfect equilibrium of the game, from a patient perspective; compute its patient payoff; and compare it to the patient payoff resulting from non-strategic behavior. We will consider only the single-state case here, for simplicity.

Proposition 4.1 (*Patient-payoff-minimizing subgame-perfect equilibrium*)

[To fill in]

Proposition 4.2 (*Payoff ratio for strategic behavior, given spending without commitment*)

[To fill in]

4.3 Spending with commitment

Let us denote the optimal spending policy given some discount rate δ by x_δ .

If a philanthropist can credibly commit to a given spending policy, then, given full knowledge of the beneficiary’s budget and discount rate, the philanthropist can implement any spending policy x for the collective budget $B_P + B$ which offers at least as much δ -discounted utility as the beneficiary could achieve by allocating B on her own. He can do this simply by demanding that he allocate

³For simplicity, we will also assume that the philanthropic risks reflected in $\delta_P(s)$ strike the philanthropist if and only if they also strike the beneficiary.

⁴For a model in which actors with different discount rates jointly set spending rates over time, see Millner and Heal (2018).

the entire collective budget according to x , and threatening to burn B_P otherwise. If B_P is high enough relative to B , this policy can achieve the patient-optimal allocation of $B_P + B$, even though the philanthropist does not own the entire collective budget.

In particular,

Proposition 4.3 (*Funding ratio needed to implement patient-optimality given commitment*)

Let x_δ denote the optimal state-dependent spending policy given discount rate $\delta(s)$, and let $v_{\delta_1, \delta_2, i}$ denote the expected $\delta_2(s)$ -discounted payoff from state s_i to following spending policy x_{δ_1} , given a budget of 1. If the philanthropist can commit to a given spending policy, he can implement the δ_P -optimal spending policy for the collective budget whenever

$$B_P \geq B \left(\left(\frac{v_{\delta_P, \delta, i}}{v_{\delta, \delta, i}} \right)^{\frac{1}{1-\eta}} - 1 \right), \quad \eta \neq 1,$$

$$B_P \geq B \left(e^{\frac{v_{\delta, \delta, i} - v_{\delta_P, \delta, i}}{[(I_{|S|} - A \circ T)^{-1} \vec{w}_\delta]_i}} - 1 \right), \quad \eta = 1.$$

This follows directly from the condition under which the patient-optimal policy for the collective budget offers the beneficiary at least as high an expected payoff as the impatient-optimal policy for the beneficiary's budget:

$$(B + B_P)^{1-\eta} v_{\delta_P, \delta, i} \geq B^{1-\eta} v_{\delta, \delta, i}, \quad \eta \neq 1, \quad (16)$$

$$v_{\delta_P, \delta, i} + [(I_{|S|} - A \circ T)^{-1} \vec{w}_\delta]_i \ln(B + B_P) \geq v_{\delta, \delta, i} + [(I_{|S|} - A \circ T)^{-1} \vec{w}_\delta]_i \ln B, \quad \eta = 1. \quad (17)$$

■

In the single-state case, we can write this threshold more simply.

Proposition 4.4 (*Funding ratio needed to implement patient-optimality given commitment, single-state case*)

In the single-state model of §2, if the philanthropist can commit to a given spending policy, he can implement the δ_P -optimal spending policy for the collective budget whenever

$$B_P \geq B \left(\left(\frac{r\eta - r + \delta_P}{r\eta - r + \delta} \right)^{\frac{\eta}{1-\eta}} - 1 \right), \quad \eta \neq 1;$$

$$B_P \geq B \left(B^{\frac{\delta_P - \delta}{\delta}} \frac{\delta^{\frac{\delta_P}{\delta}}}{\delta_P^{\frac{\delta_P}{\delta}}} e^{\frac{(r-\delta)\delta_P^2 - (r-\delta_P)\delta^2}{\delta_P \delta^2}} - 1 \right), \quad \eta = 1.$$

Again, this follows directly from the conditions under which the patient-optimal policy for the collective budget offers the beneficiary at least as high an expected payoff as the impatient-optimal policy for the beneficiary's budget:

$$\int_0^\infty e^{-\delta t} \frac{\left((B + B_P)^{\frac{r-r\eta+\delta_P}{\eta}} e^{\frac{r-\delta_P}{\eta} t} \right)^{1-\eta}}{1-\eta} dt \geq \int_0^\infty e^{-\delta t} \frac{\left(B^{\frac{r-r\eta+\delta}{\eta}} e^{\frac{r-\delta}{\eta} t} \right)^{1-\eta}}{1-\eta} dt, \quad \eta \neq 1; \quad (18)$$

$$\int_0^\infty e^{-\delta t} \ln\left((B + B_P)\delta_P e^{(r-\delta_P)t}\right) dt \geq \int_0^\infty e^{-\delta t} \ln\left(B\delta e^{(r-\delta)t}\right) dt, \quad \eta = 1.$$

■

In the $\eta = 1$ case, we can also get this expression by substituting the closed-form multi-state optimal spending policy from Proposition 3.3 into the funding ratio expression of Proposition 4.1.

4.4 Linear subsidization

Without commitment, the philanthropist can also do better if he can subsidize the beneficiary's spending, with the beneficiary maximizing her utility taking the philanthropist's subsidy policy as given. This condition might be interpreted as an assumption about the relative ease with which the patient and impatient players in this game can coordinate; in some contexts it will be more natural to refer to a continuum of uncoordinated beneficiaries. An assumption along these lines is presumably reasonable in some circumstances (such as the case of a philanthropist aiming to increase future consumption among the world's poor, discounting their future utility less than they discount their own) but unreasonable in others (such as the case of a philanthropist aiming to increase the provision of a public good whose only other provider is a single impatient philanthropist or government).

Let a_t denote the state history from 0 to t .

For simplicity, let us restrict our attention to continuous, linear subsidy policies. The philanthropist's problem is then to set an integrable function $f(a_t) \leq 1$ of effective prices, over time, for the public good under consideration.

Since we will continue to assume that the beneficiary takes this price policy as given and optimizes δ -discounted utility in the face of it, we can straightforwardly define the beneficiary's and philanthropist's budgets at time t given history a_t and price schedule f . We will denote these quantities by $B_{f,t}(a_t)$ and $B_{P,f,t}(a_t)$.

Proposition 4.5 (*The optimal price policy when taxation is feasible*)

Let $x_{\delta,\rho}(s)$ denote the optimal state-dependent spending policy given discount rate δ and price schedule $e^{-\rho t}$. If the philanthropist were able to tax expenditure in addition to subsidizing it, he would maximize δ_P -discounted utility by setting the effective price schedule

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

Furthermore, $x_{\delta,\rho}/x_{\delta_P} \triangleq x_{\delta,\rho}(s)/x_{\delta_P}(s)$ is independent of s .

Proof: First, suppose the beneficiary faces a price schedule of the form $f(t) = Ae^{-\rho t}$. She will then allocate her budget so as to achieve the post-subsidy consumption schedule that would be optimal given state-dependent discount rate $\delta(s) - \rho = \delta_P(s)$. To see this, observe that, when she does so, the expected marginal flow utility at transition time t from investing resources for expenditure at t (regardless of state) would, in the absence of the subsidy, equal $x(s)^{-\eta} e^{-\delta_P(s)t}$. In the presence of

price $Ae^{-\rho t}$, therefore, the expected marginal flow utility at t will equal $x(s)^{-\eta}e^{-\delta_P(s)t}/(Ae^{-\rho t}) = x(s)^{-\eta}e^{-\delta(s)t}/A$. The marginal utility of spending resources immediately will equal $x(s)^{-\eta}/A$. Satisfying the beneficiary's first-order conditions, the proposed spending policy will thus maximize δ -discounted utility given price schedule f .

Next, remember that, given price schedule $f(t) = e^{-\rho t}$, the beneficiary will by definition spend at proportional rate $x_{\delta,\rho}(s)$, or absolute rate $Bx_{\delta,\rho}(s)$. (Though there does not appear to be a general closed-form solution for this term, we know how to characterize it: indeed, we can straightforwardly characterize the optimal state-dependent spending policy given any price schedule $f(t)$ by multiplying the right-hand side of (11) by $f(t)^{1-\eta}$.) To induce the optimal consumption rate of $(B + B_P)x_{\delta_P}(s)$, therefore, we must have

$$A = \frac{x_{\delta,\rho}(s)}{x_{\delta_P}(s)} \left(\frac{B}{B + B_P} \right). \quad (19)$$

Finally, observe that, since price schedule \hat{f} implements the δ_P -optimal spending policy, it implements the δ_P -optimal spending policy from any point t onward, regardless of the state at t . That is, if the transition from state s_i to some s_j occurred at some time τ , the subsequent patient optimum would still be implemented by maintaining price schedule \hat{f} . We also know, however, that the subsequent patient optimum will be implemented by price schedule \hat{f} as defined at time τ . Taking a limit of the latter term, we have

$$\lim_{\tau \rightarrow 0} \frac{x_{\delta,\rho}(s_j)}{x_{\delta_P}(s_j)} \left(\frac{B_\tau}{B_\tau + B_{P,\tau}} \right) e^{-\rho(t-\tau)} = \frac{x_{\delta,\rho}(s_i)}{x_{\delta_P}(s_i)} \left(\frac{B}{B + B_P} \right) e^{-\rho t}. \quad (20)$$

It follows that $x_{\delta,\rho}(s_i)/x_{\delta_P}(s_i) = x_{\delta,\rho}(s_j)/x_{\delta_P}(s_j)$ for arbitrary i, j , and thus that $x_{\delta,\rho}(s)/x_{\delta_P}(s)$ is independent of s . ■

As usual, in the single-state case, we can describe this policy more precisely.

Proposition 4.6 (*The optimal price schedule when taxation is feasible, single-state case*)

Given a single state, 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 $f(t)$, she allocates her budget according to allocation $y_f(t)$ such that, for some constant k ,

$$\frac{\partial}{\partial[y_f(t)]} \left[e^{-\delta t} \frac{\left(\frac{e^{rt} y_f(t)}{f(t)} \right)^{1-\eta}}{1-\eta} \right] = k \quad \forall t. \quad (21)$$

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

$$y_{\hat{f}}(t) = k^{\frac{-1}{\eta}} e^{\frac{r-r\eta+\delta_P\eta-\delta\eta-\delta_P}{\eta}t} \left(\frac{B}{B + B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}}. \quad (22)$$

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

$$k^{\frac{-1}{\eta}} \left(\frac{B}{B+B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}} \int_0^\infty e^{\frac{r-r\eta+\delta_P\eta-\delta\eta-\delta_P}{\eta}t} dt = B; \quad (23)$$

$$k^{\frac{-1}{\eta}} \left(\frac{B}{B+B_P} \right)^{\frac{\eta-1}{\eta}} \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{r\eta - r + \delta_P} \right)^{\frac{\eta-1}{\eta}} = B \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta} \right). \quad (24)$$

Substituting (24) into (22) and remembering that $x(t) = e^{rt}y(t)$, we see that the beneficiary now spends according to schedule

$$x_{\hat{f}}(t) = B \left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta} \right) e^{\frac{r+\delta_P\eta-\delta\eta-\delta_P}{\eta}t}. \quad (25)$$

At each t , the beneficiary thus consumes

$$\frac{x_{\hat{f}}(t)}{\hat{f}(t)} = (B+B_P) \left(\frac{r\eta - r + \delta_P}{\eta} \right) e^{\frac{r-\delta_P}{\eta}t}. \quad (26)$$

From Proposition 2.1, 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. \blacksquare

In short, Proposition 4.6 finds that $(r\eta - r - \delta_P\eta + \delta\eta + \delta_P)/(r\eta - r + \delta_P)$ is the closed-form value of Proposition 4.3's $x_{\delta,\rho}/x_{\delta_P}$ in the single-state case.

As we can see, a price schedule of the form $Ae^{-\rho t}$ lowers the beneficiary's effective discount rate by ρ across reasonably general circumstances. This result stands in contrast to the fact, well-known in the optimal taxation literature, that policymakers cannot generally implement their first-best tax and subsidy policies over time without the power to commit (see e.g. Klein and Ríos-Rull, 2003; Benhabib and Rustichini, 1997).

Note, however, that this result relies on the feature of our setup that the utility achieved by spending at a time t accrues at t . It would not hold in a more general setting, such as an arbitrary Markov decision process. If a given state held the fleeting opportunity to fund a project most of whose value would accrue in the future, the patient might well try to induce a faster proportional spending rate than the impatient.

Propositions 4.3 and 4.4 may offer useful advice for a patient policymaker in the setting described. 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 his resources until $t^* : \hat{f}(t^*) = 1$, setting $f(t) = 1$ for $t < t^*$ and $f(t) = \hat{f}(t)$ for $t \geq t^*$.

Corollary 4.7 (*Funding ratios needed to implement patient-optimality without commitment*)

Without the ability to credibly commit, the philanthropist can implement the δ_P -optimal spending policy with a linear subsidy policy so long as

$$\frac{B_P}{B} \geq \frac{x_{\delta,\rho}}{x_{\delta_P}} - 1.$$

In the single-state case, the condition can be written

$$\frac{B_P}{B} \geq \frac{\eta\rho}{r\eta - r + \delta_P}.$$

This follows immediately from setting $\hat{f}(t) = 1$ in Propositions 4.3 and 4.4 and solving for B_P/B .

Comparing these terms to the terms in Propositions 4.1 and 4.2 sheds light on the benefits of a mechanism for credible commitment, and might lay the groundwork for an analysis of how much patient philanthropists should be willing to pay for one. But we will not pursue this question further here.

Proposition 4.8 (*Markov perfect equilibrium*)

Consider the strategy profile in which the philanthropist invests all his resources until $B_P/B = x_{\delta,\rho}/x_{\delta_P} - 1$ and then subsidizes according to price policy \hat{f} , and the beneficiary spends so as to maximize δ -discounted utility under the assumption that the philanthropist will do the above. With a_t as our state variable, this strategy profile is in Markov perfect equilibrium.

Proof: Consider a discrete-time analog of the continuous-time game presented here. We begin in state s_i . Every θ units of time, the philanthropist sets a price $f(t)$ ($t = 0, \theta, 2\theta, \dots$), and the beneficiary subsequently chooses an annual expenditure rate $x(t) \in [0, \min(B_t, B_{P,t}/(1 - f(t)))/\theta]$. The state then changes, perhaps, with $T_{i,i} = \sum_{j \neq i} T_{i,j}$. All resources not consumed at a stage are invested at interest rate $r(s(t))$, such that

$$B_{t+\theta} = (B_t - \theta x(t))(1 + r(s(t)))^\theta, \quad (27)$$

$$B_{P,t+\theta} = (B_{P,t} - (1/f(t) - 1)\theta x(t))(1 + r(s(t)))^\theta.$$

In particular, consider the strategy profile in which the philanthropist sets $f(t) = 1$ so long as $B_{P,t}/B_t < x_{\delta,\rho}/x_{\delta_P} - 1$, where $x_{\delta,\rho}$ denotes the optimal spending policy in discrete time given discount rate δ and price policy $f(t) = (1 - \rho)^{\theta t}$, as characterized by multiplying the right-hand side of Proposition 3.1 by $(1 - \rho)^\theta$, and sets

$$f(a_t) = \frac{x_{\delta,\rho}}{x_{\delta_P}} \left(\frac{B_{t^*}}{B_{t^*} + B_{P,t^*}} \right) (1 - \rho)^{\theta t}, \frac{B_{P,t}}{B_t} \geq \frac{x_{\delta,\rho}}{x_{\delta_P}} - 1, \quad (28)$$

where t^* is the time at which the term at right reaches equality; and in which the beneficiary spends so as to maximize δ -discounted utility given this price policy.

The beneficiary's strategy trivially constitutes a best response to the philanthropist's, since as specified at the beginning of this section, the beneficiary acts as a price-taker, maximize her utility taking the subsidy policy as given.

It follows from Proposition 4.3, as rewritten in discrete time, that from any history $a_t : B_{P,f,t}(a_t)/B_{f,t}(a_t) \geq x_{\delta,\rho}/x_{\delta_P} - 1$, price policy (25) implements the δ_P -optimal subsequent allocation of budget $B_t + B_{P,t}$. It follows that no deviation from such a state could increase the philanthropist's payoff.

Finally, consider a history $a_t : B_{P,t}/B_t < x_{\delta,\rho}/x_{\delta_P} - 1$, and consider a one-shot deviation by the philanthropist in which he sets $f(a_t) < 1$.

This deviation lowers δ_P -discounted utility. To see this, observe that $B_{P,\theta}/B_\theta$ is increasing in $f(0)$. On setting $f(0) < 1$, therefore, at least as many periods will have to elapse before the condition for proportional subsidies to begin is met; and if equally many periods are met, the subsidy rate will nevertheless be lower at each period.

Let us say that subsidization would have begun at t^* and instead begins at $t^{**} \geq t^*$. The deviation will then proportionally increase prices at all $t \geq t^{**}$, and will monotonically increase prices over $\{t^*, \dots, t^{**} - \theta\}$. By the homotheticity of the isoelastic utility function, this monotonic transformation of the price schedule will monotonically shift the allocation of the collective budget at $t = \theta$ from later to earlier periods. Furthermore, the deviation unambiguously increases spending at $t = 0$, decreasing the absolute budget $B_\theta + B_{P,\theta}$ to be allocated subsequently. The deviation thus induces a strict monotonic shift in the budget allocation at $t = 0$ from later to earlier periods.

If the philanthropist waits until t^* to subsidize, the marginal δ_P -discounted utility to allocating resources for investment until t is decreasing until t^* and flat for $t \geq t^*$. The budget reallocation induced by the deviation from this policy thus monotonically shifts resources from periods where they produce more δ_P -discounted utility to periods in which they produce less, lowering total δ_P -discounted utility. By the one-shot deviation principle, it follows that this strategy profile is in equilibrium.

Finally, the continuous-time game above is the limit of this discrete-time game as $\theta \rightarrow 0$, and the strategy profile under consideration consists of pure strategies. It follows (Simon and Stinchcombe, 1989) that the corresponding pure strategies of the continuous-time game are in equilibrium. ■

In the single-state case, we can examine this result more precisely.

Proposition 4.9 (*The impatient response to the patient philanthropist, single-state case*)

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

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

and subsidizes at prices given by $f(t) = \hat{f}(t - t^*)$ 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

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

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

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

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 her consumption so as to maximize δ -discounted utility, she will be indifferent between allocating marginal consumption to $t = 0$ and to any $t \in [0, t^*]$. That is, as in (4),

$$x_f(0)^{-\eta} = e^{(r-\delta)t} (x_f(t))^{-\eta}; \quad (29)$$

$$x_f(t) = x_f(0) e^{\frac{r-\delta}{\eta} t}. \quad (30)$$

Let us now define

$$D \triangleq \frac{x_f(0)\eta}{B(r\eta - r + \delta)}, \quad (31)$$

so that

$$x_f(t) = DB \left(\frac{r\eta - r + \delta}{\eta} \right) e^{\frac{r-\delta}{\eta} t}, \quad t \leq t^*. \quad (32)$$

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

$$\left(DB \left(\frac{r\eta - r + \delta}{\eta} \right) \right)^{-\eta}. \quad (33)$$

Now, observe that 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, her endowment is of size

$$B_t = B e^{(r - D \left(\frac{r\eta - r + \delta}{\eta} \right)) t}. \quad (34)$$

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

$$t^* : \frac{B_P e^{rt^*}}{B e^{(r - D \left(\frac{r\eta - r + \delta}{\eta} \right)) t^*}} = \frac{\eta(\delta - \delta_P)}{r\eta - r + \delta_P}; \quad (35)$$

$$t^* = \frac{1}{D} \ln \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right) \frac{\eta}{r\eta - r + \delta}. \quad (36)$$

Substituting (33) into (29), we have the beneficiary's spending rate at t^* :

$$x_f(t^*) = DB \left(\frac{r\eta - r + \delta}{\eta} \right) \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right)^{\frac{r-\delta}{D(r\eta - r + \delta)}}. \quad (37)$$

Substituting (33) into (31), we also have

$$B_{t^*} = B \left(\eta \frac{B}{B_P} \frac{\delta - \delta_P}{r\eta - r + \delta_P} \right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1}. \quad (38)$$

Let us substitute (35) for B into (22) and so re-index t^* to 0. Since the philanthropist begins subsidizing to achieve price schedule \hat{f} for $t \geq t^*$, we have

$$x_f(t) = B\left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta}\right)\left(\eta\frac{B}{B_P}\frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1} e^{\frac{r + \delta_P\eta - \delta\eta - \delta_P}{\eta}(t - t^*)}, t \geq t^*. \quad (39)$$

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

$$x_f(t^*) = B\left(\frac{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}{\eta}\right)\left(\eta\frac{B}{B_P}\frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)^{\frac{r\eta}{D(r\eta - r + \delta)} - 1}. \quad (40)$$

Setting (37) equal to (32) and solving for D , we have

$$D = \frac{\ln\left(\eta\frac{B}{B_P}\frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)}{W\left(\eta\frac{B}{B_P}\frac{\delta - \delta_P}{r\eta - r + \delta_P}\frac{r\eta - r + \delta}{r\eta - r - \delta_P\eta + \delta\eta + \delta_P}\ln\left(\eta\frac{B}{B_P}\frac{\delta - \delta_P}{r\eta - r + \delta_P}\right)\right)}. \quad (41)$$

Terms (29), (36), (33), and (38) constitute our result. ■

Proposition 4.10 (*Optimality 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 :

$$\frac{\partial}{\partial y(t)} \left[e^{-\delta t} \ln\left(\frac{e^{rt}y(t)}{f(t)}\right) \right] = \frac{\partial}{\partial y(t)} \left[e^{-\delta t} \ln(e^{rt}y(t)) \right] = \frac{e^{-\delta t}}{y(t)}. \quad (42)$$

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 his endowment directly to augmenting consumption at the time-periods of his choosing, without worrying about substitution on the part of the beneficiary.

When the philanthropist follows price schedule

$$\hat{f}(t) = e^{(\delta_P - \delta)t} \frac{B_{t^*}}{B_{t^*} + B_{P,t^*}} \frac{\delta}{\delta_P}, \quad t \geq t^* = \ln\left(\frac{B(\delta - \delta_P)}{B_P\delta_P}\right)/\delta, \quad (43)$$

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

$$\frac{\partial}{\partial [y(t)/\hat{f}(t)]} \left[e^{-\delta_P t} \ln\left(\frac{e^{rt}y(t)}{\hat{f}(t)}\right) \right] = \frac{1}{(B_{t^*} + B_{P,t^*})\delta_P}, \quad (44)$$

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 $t \geq t^*$. (Recall that, by (42), neither the taxes nor the subsidies would be distortionary.)

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

Proposition 4.11 (*Payoff ratio for strategic behavior, given linear subsidization*)

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.

We can prove this analytically where $\eta = 1$. The result for arbitrary η is given by Mathematica.⁵

Proof: The payoff from following f^* is

$$\begin{aligned} & \int_{t^*}^{\infty} e^{-\delta_P t} \left(\ln \left(\frac{1}{f^*(t)} B \delta^{(r-\delta)t} \right) - \ln \left(B \delta^{(r-\delta)t} \right) \right) dt \\ &= \int_{t^*}^{\infty} e^{-\delta_P t} \ln \left(\frac{1}{f^*(t)} \right) dt \\ &= \frac{\delta - \delta_P}{\delta_P^2} \left(\frac{B_P}{B} \frac{\delta_P}{\delta - \delta_P} \right)^{\frac{\delta_P}{\delta}}. \end{aligned} \tag{45}$$

The payoff from an immediate transfer is

$$\begin{aligned} & \int_0^{\infty} e^{-\delta_P t} \left(\ln \left((B_P + B) \delta e^{(r-\delta)t} \right) - \ln \left(B \delta e^{(r-\delta)t} \right) \right) dt \\ &= \int_0^{\infty} e^{-\delta_P t} \ln \left(\frac{B_P}{B} + 1 \right) dt \\ &= \frac{\ln \left(\frac{B_P}{B} + 1 \right)}{\delta_P}. \end{aligned} \tag{46}$$

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

$$\lim_{\frac{B_P}{B} \rightarrow 0} \frac{\frac{\delta - \delta_P}{\delta_P} \left(\frac{B_P}{B} \frac{\delta_P}{\delta - \delta_P} \right)^{\frac{\delta_P}{\delta}}}{\ln \left(\frac{B_P}{B} + 1 \right)} = \infty, \tag{47}$$

as desired. \blacksquare

The overarching conclusion of the shallow investigation above is that, in the presence of impatient funding, the patient generally ought to do something like the following. They should try to invest

⁵[Put code in an online appendix]

until the impatient are spending their private budgets at what, from a patient perspective, would be the optimal rate for the collective budget, and then disburse their funds so as to fill the gap between the spending schedule of the impatient and the patient-optimal subsequent allocation of the collective budget. Where it is possible to implement a disbursement plan along these lines, the payoff to doing so can be much larger than the payoff to allowing the global spending schedule to be set impatiently.

4.5 Examples

Let us first consider the case of patient philanthropists aiming to do good by increasing global human consumption-based utility—the world’s largest cause area by far.

Suppose that $\eta = 1$, $r = 4\%$, $\delta = 1.5\%$, $\delta_P = 0.5\%$,⁶ $B = \$3.4\text{Q}$,⁷ and $B_P = \$20\text{B}$.⁸ Then, by Propositions 4.4 and 4.8, the philanthropist maximizes δ_P -discounted utility by investing until year $t^* = 849$ and then subsidizing consumption according to price-schedule \hat{f} . Doing so will produce consumption increases after year 849 (i.e. calendar year 2868) of

$$\begin{aligned} & \left(\frac{1}{\hat{f}(t-849)} - 1 \right) x(t) \\ &= \left(e^{(\delta-\delta_P)(t-849)} \frac{(B_{849} + B_{P,849})\delta_P}{B_{849}\delta} - 1 \right) B_{849}\delta e^{(r-\delta)(t-849)} \\ &= \left(e^{0.01(t-849)} \frac{(3.4\text{Q} \cdot e^{0.025 \cdot 849} + 20\text{B} \cdot e^{0.04 \cdot 849})0.005}{3.4\text{Q} \cdot e^{0.025 \cdot 849} \cdot 0.015} - 1 \right) 3.4\text{Q} \cdot e^{0.025 \cdot 849} \cdot 0.015 e^{0.025(t-849)} \\ &= 8.42 \cdot 10^{22} (e^{0.035(t-849)} - e^{0.025(t-849)}) \end{aligned} \tag{48}$$

for an expected δ_P -discounted payoff of

$$\int_{849}^{\infty} e^{-0.005t} \left(\ln(8.42 \cdot 10^{22} \cdot e^{0.035(t-849)}) - \ln(8.42 \cdot 10^{22} \cdot e^{0.025(t-849)}) \right) dt = 5.73. \tag{49}$$

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 a payoff of

$$\int_0^{\infty} e^{-0.005t} (\ln((3.4\text{Q} + 20\text{B}) \cdot 0.015 e^{0.025t}) - \ln(3.4\text{Q} \cdot 0.015 e^{0.025t})) dt = 0.0012, \tag{50}$$

for a payoff ratio of $5.73/0.0012 = 4,775$.

⁶This is the fully patient rate if the expected lifespan of a fund designed for the long term is about 200 years, or if its half-life is about 137 years.

⁷Global output in 2018 was \$86T (World Bank, 2019). If this grows at the growth rate, the present value of future output, after discounting by the interest rate, is $\$86\text{T}/(0.04 - 0.015) = 3.4\text{Q}$. Note that this is substantially larger than “world wealth” as typically defined (e.g. by Credit Suisse Research Institute, 2018), primarily because the latter excludes all present and future human capital.

⁸This is intended to be a rough estimate of the resources possessed by the Effective Altruism community, including discounted future wages.

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 maximally cost-effective intervention for the sake to the world's poorest today. If $\eta = 1.2$, transfers to the world's poorest, making \$300 per year, produce 67x more welfare than transfers to those earning the global median income of \$10,000 per year; and the most cost-effective poverty alleviation efforts today are estimated to be up to 61x more cost-effective than cash transfers (GiveWell, 2019). As 4,775 is higher even than the product of these estimates, however, these figures suggest that patience ("getting the time right") is worthwhile even if comes at the cost of *all* targeting ("getting the place right").

Note also that the 849-year waiting period is many times 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—and as one can calculate from the multi-state case—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 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

$$9.550 \cdot 10^8 \cdot e^{0.035(t-46)} - 9.475 \cdot 10^8 \cdot e^{0.025(t-46)} \quad (51)$$

for an expected δ_P -discounted payoff of

$$\int_{46}^{\infty} e^{-0.005t} \left(\ln(9.550 \cdot 10^8 \cdot e^{0.035(t-46)}) - \ln(9.475 \cdot 10^8 \cdot e^{0.025(t-46)}) \right) dt = 319. \quad (52)$$

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

$$\int_0^{\infty} e^{-0.005t} (\ln(40B \cdot 0.015e^{0.025t}) - \ln(20B \cdot 0.015e^{0.025t})) dt = 138, \quad (53)$$

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

As we can see, the proportional benefits to acting strategically 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.

5 Applications to philanthropic cause-areas

5.1 Direct efforts to increase near-term human welfare

5.1.1 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. Policymakers 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.
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.⁹

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

On the relationship between the interest and social discount rates As defined above, r is the rate at which individuals discount consumption. Observe that, under ordinary circumstances, r will equal the market interest rate. At least when evaluating projects whose payoffs’ time-profiles

⁹Note that, under these conditions, η simultaneously equals (a) individuals’ coefficients of relative risk aversion, (b) the inverse of individuals’ elasticities of intertemporal substitution, and (c) the degree of social inequality aversion in consumption.

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 whose (effective) time preference rate is $\delta_P < \delta$ —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.

Empirical estimates For the rest of §5.1, I will assume values of $\delta = 0.02$, $g = 0.02$, and $\eta = 1.2$. For simplicity, I will also assume that $\delta_P = 0$, but it should be clear that identical reasoning applies to any case in which $\delta_P < \delta$.

Note that optimal disbursement schedules can be highly sensitive to our choice of η , and that unfortunately, estimates of η in particular vary widely.

Research 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, 2006). Global survey data on the relationship between consumption and self-reported wellbeing suggests the same (Gandelman and Hernández-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.

On the other hand, 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. These estimates are quite extreme; if $\eta = 5$, for example, marginal consumption for someone consuming \$30,000 per year produces over four hundred times more welfare than marginal consumption for someone consuming \$100,000 per year. This anomalous financial data, however, may be 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.

In any event, for now we will use $\eta = 1.2$, bearing in mind that this is a relevant source of uncertainty.

5.1.2 A common mistake: “Discount by $\eta g^* > \delta + \eta g$ ”

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 at least in the domain of global poverty alleviation, 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.¹⁰ Letting g^* denote the consumption growth rate among the world’s poorest, this amounts roughly to the claim that $\eta g^* > \delta + \eta g$.¹¹ 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 any further considerations.

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 a mistake in the process which biased his analysis in favor of giving later,¹² and omitting many of the timing considerations discussed here. 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

¹⁰Alexander (2013) reports Elie Hassenfeld, a cofounder of EA charity evaluator GiveWell, making essentially this point.

¹¹Note 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 claim 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.

¹²Dickens assumes logarithmic utility. (Recall that logarithmic utility is the special case of isoelastic utility with $\eta = 1$.) 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. And some estimates of η are substantially higher than 1, as noted in §5.1.1.

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).

5.1.3 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 now than by giving in the immediate future.

We cannot, however, naively extrapolate this instantaneous relationship. When considering whether to give now or in the more distant future, 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 .

To see this, observe first 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.

More concretely, suppose that the developing world must grow $n \geq 1$ times as wealthy before it gain functioning access to the global credit market, and suppose that it is currently growing $m \geq 1$ times as fast as would be necessary to maintain the condition that giving now is at least weakly preferable to giving in the immediate future: that is, that $g^* = \frac{mr}{\eta}$. Then it will gain access to the global credit market in $\ln(n) \frac{\eta}{mr}$ years, and subsequently grow at rate g . Resources invested at rate r will then have grown by more than the cost of buying a unit of welfare in the developing world at

$$t : e^{\ln(n) \frac{\eta}{mr} \cdot \eta \cdot \frac{mr}{\eta} + (r - \delta)(t - \ln(n) \frac{\eta}{mr})} = e^{rt}, \quad (54)$$

$$t = \ln(n) \eta \left(1 - \frac{r - \delta}{mr}\right) / \delta,$$

for a lower bound (varying m) of $\ln(n) \eta (1 - r + \delta) / \delta$ and an upper bound of $\ln(n) \eta / \delta$. If $\delta = 0.02$, $\eta = 1.2$, and $n = 30$ —that is, if the world’s poorest countries will have functioning credit markets only once their incomes reach roughly that currently enjoyed by Poland (CITE)—then the time range in question is 93–204 years.

Furthermore, even if, for the sake of argument, the world's poorest countries grew forever as closed economies at some rate $g^* \gg g$, the rate of increase in the cost of producing a unit of welfare as efficiently as possible still could not permanently exceed ηg . If it did, 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 starting endowment. In other words, in the long run, we should discount by something no more than 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 $\eta = 1.2$, 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 $1.044^{279} = 165,000$, and this effect swamps the satiation effect, as we have seen.

The argument from flow-through effects fails for similar reasons. Educating a child, or curing her of some debilitating disease, may well increase her future earnings by more than the invested cost of the treatment over a similar timeframe. But she will not then invest the entirety of the fruits of her labor at any rate greater than or equal to r . Instead, as derived in Proposition 2.1, she will split between consumption and (presumably local) investment so that her assets grow at rate $(r - \delta)/\eta$ —which is therefore, of course, the local economic growth rate.¹³

Christiano (2013) offers a similar reason for doubting that the rate γ at which a philanthropic intervention increases consumption growth could, in the long run, nontrivially exceed g . This is the observation that if it did, the flow-through effects of the intervention would eventually be larger than the rest of the world economy, and individual historical acts of effective philanthropy would have single-handedly pushed up the global growth rate to γ .

In short, even under what appear to be the most favorable plausible assumptions, an intervention to increase the productivity of the poor should roughly be expected to offer temporary returns of $\gamma > g^*$, followed by medium-term returns of $g^* > g$, followed by long-run returns of g . The plan

¹³We can also find the $(r - \delta)/\eta = g$ relationship by rearranging the Ramsey formula.

to maximize expected welfare, at least according to the line of reasoning given in this paper, will therefore resemble that sketched in §4.4: investment for several centuries, followed by an exponential disbursement schedule. (If we do not believe that individuals, governments, and other philanthropists will respond to such a plan by slightly lowering their rates of investment, we can implement this straightforwardly. If we believe that they will, the plan would have to involve partnering with future governments to implement lower consumption taxes, or consumption subsidies, in the event that the fund survives as planned, in which case it will have grown to a substantial proportion of global wealth.) If we believe that the developing world is currently undergoing permanent and inevitable catch-up growth at rate $g^* > g$, and that the catch-up process will take about T years in expectation, we can loosely model this by positing two states s_1, s_2 with $h(s_1) = e^{\eta(g^* - g)T}$, $h(s_2) = 1$, $T_{1,2} = 1/T$, and $T_{2,1} = 0$, and, as before, conceiving of the world’s self-interested, impatient population as a very wealthy continuum of impatient philanthropists. If we have more fine-grained beliefs about short-term trends and fluctuations in the cost of doing good in the developing world, we can at least loosely model these as well. But as the estimates above suggest, these tweaks are unlikely to overturn the overarching conclusion.

Finally, note that this argument—that the consumption of the world’s poorest cannot grow forever at a rate g^* which exceeds the average consumption growth rate g —only establishes ηg as an *upper bound* on the long-run rate at which it is growing more costly to augment human welfare. The world’s poorest (as a quantile) have historically witnessed slower growth than the world as a whole: they lived at subsistence now, and they lived at subsistence in the past. When we remember the possibility that the poorest will continue to get richer at rates below rather than above g —including the possibility that some parts of the world will remain in, or return to, Malthusian equilibria in the very long run—investing to help the future poor appears even more appealing.

5.2 Longtermist efforts

It is also commonly argued in the EA community that philanthropists can do more good through efforts aimed directly at improving the expected value of the distant future, such as existential risk reduction, than through poverty alleviation (Bostrom, 2003; Greaves and MacAskill, 2019).

The claim that the present is a time at which we are uniquely well-poised to improve the expected value of the distant future was put perhaps most famously by Parfit (2011), who hypothesized that he was living at the “hinge of history”. The idea 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 the future, 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 might say that a state is “hingey” in proportion to the maximum expected good that can be done in that state per unit of philanthropic resources. 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.

More formally still, though somewhat less generally, we might construct a variable $v(t)$ to capture the all-things-considered expected value of the future. (This would be what Bostrom (2014) calls an “evaluation heuristic”, or what I like to call a “longtermist economic index” (Greaves et al., 2019).) If patient philanthropists have the opportunity to push up $v(t)$ at a rate isoelastic in their spending rate $x(t)$, hingeyness would then be captured in the model of §3 as the scale parameter $h(s)$. Letting s^* denote the current state, the “hinge of history hypothesis” is then the hypothesis that $h(s^*) \gg h(s) \forall s \neq s^*$, and that $\sigma(s^*) \gg 0$.

The benefits to acting strategically in funding longtermist efforts is likely smaller than that to acting strategically in funding more direct efforts to increase human welfare, since the latter are already much better funded. But there are still some such benefits. The U.S. Congress, for example, is willing to allocate NASA enough resources to monitor only approximately 90% of potentially earth-bound asteroids (Nordhaus, 2011; NASA, 2007). A patient philanthropist who wishes to invest his resources for the sake of more thorough asteroid monitoring in the future runs the risk of inducing Congress to spend their monitoring budget more quickly. Patient philanthropists might thus consider engaging in public-private partnerships in which they match future asteroid monitoring expenditures according to some exponential schedule.

Let us now examine the implications of the hinge of history hypothesis more closely. Parfit did not mean that the hingeyest year was literally 2011, of course. In fact he goes on to clarify his position, saying that he only believes that it is of paramount importance that humanity “act wisely in the next few centuries”. If we interpret this to mean that we currently face, say, a 0.3% annual probability of existential catastrophe, then his warning should encourage us to shift our focus to lowering this risk, but it has relatively small implications for the timing problem. It only increases the philanthropic discount rate from, say, 0.5% to 0.8%. That is, in the event that the rate of existential risk reduction is roughly logarithmic in the spending rate, it only increases the optimal spending rate from roughly 0.5% to 0.8%. On this view, the “hinge of history” essentially constitutes a centuries-long plateau on which philanthropists should try to spread their spending relatively thinly.

But perhaps it is important to be more precise about whether the times are currently getting more important or less. If 2011 was the hingeyest year, the patient philanthropist should currently spend his resources on high-impact opportunities quickly, since such opportunities are running out. On the other hand, consider the possibility that the most important year in this sense is 2061: that this is the year after World War III, say, during which lobbyists and policymakers will write the constitution that rules the world forever after. In this scenario, 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.¹⁴

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?

In short, if we really are living at the most important era in history by a large enough margin, and if this era is sure enough to be very brief, then, trivially, philanthropic resources ought to be spent quickly. Outside of such extreme circumstances, however, the magnitude of “hingeyness” considerations on the optimal spending rate is slighter than it might first appear; and even the sign of “hingeyness” considerations on the current optimal spending rate appears to be ambiguous, even on the “hinge of history hypothesis”, which should render current spending most appealing.¹⁵

Here is another way to express the thought. 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

¹⁴This model of philanthropy is similar to that recommended by Hanson (2018a), who suggests that patient philanthropists “fund a long-lived organization that invests and saves its assets, and then later spends those assets to influence some side in a fight”.

¹⁵Interested readers can test the consequences of their own beliefs about the hingeyness process using the online tool accompanying this paper: <https://www.philiptrammell.com/dpptool/>.

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.

6 Further considerations

6.1 Relationships to endogenize

The limitations of the timing model above are that η and T 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 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.

6.1.1 Endogenous learning

Since learning enters the picture here only via h , our choice to make T exogenous also appears to render learning exogenous. In the special case that research proportionally increases the cost-effectiveness of all future spending, it is functionally equivalent to investment, so the current model does let us determine how much to spend on it. But a more thorough account of endogenous learning, and intervention reversibility, are conspicuously absent from the picture so far—especially given how centrally they feature in other models of optimal timing, e.g. Pindyck (2002).

6.1.2 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 $\frac{r-\delta_P}{\eta} > g$. 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 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.”)
- [I should look into “evolutionary finance” here.]

6.1.3 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 as a proportion of global wealth. This relationship may be positive or negative. People might grow more inclined to seize it, for example, or it might grow better able to defend itself.

Gwern (2012) argues that the high institutional mortality rate (estimated at 99% per century), and the rarity of long-lasting institutions, demonstrates that the expropriation rate is high and that this constitutes a definitive argument against attempting long-term saving. This argument appears to be predicated on at least three misunderstandings.

First, while institutions in general have a high mortality rate, the relevant mortality rate is of course that among the much smaller group of institutions designed to be long-lasting. Furthermore, though many ancient institutions were doubtless intended to be long-lived, it is not clear *how much* they valued their longevity. If we put, say, a decade and a hundred million dollars into asking the world’s best strategists how to design a maximally resilient and stable institution—which the EA community could easily afford to do—we will be in a very small historical reference class.

Second, it is misleading to cite the large numbers of failed philanthropic institutions (such as Islamic waqfs) which were intended to be permanent, since their closures were not independent. For illustration, if a wave of expropriation (say, through a regional conquest) is a Poisson process with $\lambda = 0.005$, then the probability of a thousand-year waqf is 0.7%. Splitting a billion-dollar waqf into a billion one-dollar waqfs, and observing that none survive the millennium, will give the impression that “the long-term waqf survival rate is less than one in one billion”.

Third, and most importantly, no part of the case for investing requires that δ_P be low. However high it is, investors will generally be compensated for it with a higher interest rate, since $r = \delta_P + \rho + \eta g$. A historical case against long-term investing thus requires a demonstration that the expropriation rate *grows with fund size*, and in particular that there is a fund size \bar{B} such that $\delta_P(B) > \delta \forall B \geq \bar{B}$. This claim is sometimes made, but so far, to my knowledge, has not been justified. Furthermore, I suspect that it is false, for three reasons.

- There have certainly been historical episodes in which institutions have been expropriated because of their relative wealth, such as King Henry VIII’s expropriation of England’s Catholic

monasteries. But I suspect that there are even more episodes (consider the pogroms) in which institutions have been expropriated because their relative poverty left them defenseless.

- Even though only a relatively small proportion of large institutions are ancient, it seems clear that an even smaller proportion of small institutions are ancient.
- Except in the rare event of a truly popular revolution, institutions are only expropriated by wealthier institutions. That is, the governments that have historically expropriated philanthropic institutions are themselves institutions that invested their assets (albeit often barbarically) until they had grown from households to fiefdoms to nations. Even if the expropriation rate increases somewhat for institutions just big enough to threaten or tempt the governments nearby them, therefore, it seems this risk must be offset by the non-negligible reward that awaits an institution big enough to achieve some autonomy and start doing the expropriating.

A handful of relatively recent attempts explicitly to found long-term trusts have met with with partial success (Benjamin Franklin) or comical failure (James Holdeen). Unfortunately, there have not been enough of these cases to draw any compelling conclusions. In any event, more research on the historical relationship between δ_P and fund size would be valuable.

6.1.4 Uncertainty about long-run r , δ_P , and η

[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 to change.

6.1.5 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 grand 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 $\eta = 2$. [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.

6.2 Investment strategies

6.2.1 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).
- To the extent that the “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.

6.2.2 Movement building

[Informal notes:] It was noted in §5 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.

6.2.3 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 natural way to account for this possibility, I believe, is to model “the world at large” as a risky asset which every patient philanthropst already possesses. 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 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.

6.2.4 Leverage

[Informal notes:] In the face of risky investment opportunities, a cause’s patient and impatient funders may differ in their willingness to take on leverage. Accounting for this may produce game-theoretic complications beyond those present in the case of risk-free investment, as explored in §4.

7 Conclusion

If our aim as philanthropists is simply to do the most good, 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 §6.2. 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 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.

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, as a generalized consumption-smoothing model reveals, there are at least three clear arguments for thinking that, at least in ordinary circumstances, patient philanthropists should prefer waiting.

First: Most individuals face mortality risks, and rates of pure time preference, that do not appear in the patient philanthropist’s utility function. Patient philanthropists thus have much stronger preferences for what happens to the distant future than self-interested actors do—stronger, that is, than those of almost 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: consider, by illustration, that we can 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, the patient can give the impatient their assets in the short term in exchange for more their assets in the long term—or in a setting without lending, take investment opportunities at lower interest rates than the impatient demand. When a patient philanthropist’s objectives are *orthogonal* to those of the world’s impatient actors, therefore, both parties should agree to an exchange in which “the impatient receive the near-term” in exchange for “the patient inheriting the long term”.

Second: 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. In other words, 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.

Third: A patient philanthropist’s objectives will often *overlap* with those of less patient funders. Since these objectives presumably face diminishing returns to spending, and since the impatient are inclined to spend their endowments in the near-term anyway, this overlap gives the patient philanthropist yet more reason to delay spending.

As noted in §6, the generalized consumption-smoothing model explored here leaves many questions open. At least some of the considerations still unexplored, however, further strengthen the case for philanthropic investment. Consider, for example, the “competitive” possibility that one’s philanthropic impact will ultimately depend not on how much one owns but on what fraction of the world’s resources one owns. In this scenario, investing is extremely important, since—as noted for example by Hanson (2018c)—patience cannot remain unusual forever. The patient will generally invest more than the impatient, and as they do, they will come to own more of the world. Once most of the world belongs to the patient, resources will subsequently be allocated roughly along patient lines, and the opportunity to “cheaply 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 in a competitive scenario, we know at least that we are living at a special time with respect to the value of saving.

The incomplete model developed here suggests that many philanthropists could fulfill their own

values substantially more effectively by spending more slowly. It is not clear whether further modeling would weaken or strengthen this suggestion, but we can at least conclude from our model so far that some arguments given in favor of giving now are mistaken. In short, patient philanthropists should reconsider their decision to spend as quickly as they do—and philanthropically-minded economists should consider that the problem of discounting for patient philanthropists is important, decision-relevant, and amenable to further fruitful exploration.

8 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] Belfiori, M.E. (2017). “Carbon Pricing, Carbon Sequestration, and Social Discounting”, *European Economic Review* 96: 1–17.
- [4] Bernheim, B.D. (1989). “Intergenerational Altruism, Dynastic Equilibria and Social Welfare”, *Review of Economic Studies* 56 (1): 119–28.
- [5] Benhabib, J., and A. Rustichini (1997). “Optimal Taxes without Commitment”, *Journal of Economic Theory* 77(2): 231–59.
- [6] Bergstrom, T., L. Blume, and H. Varian (1986). “On the Private Provision of Public Goods”. *Journal of Public Economics* 29(1): 25–49.
- [7] Bostrom, N. (2003). “Astronomical Waste: The Opportunity Cost of Delayed Technological Development”, *Utilitas* 15(3): 308–14.
- [8] Bostrom, N. (2014). “Crucial Considerations and Wise Philanthropy”. URL: <http://www.stafforini.com/blog/bostrom/> (accessed 1 Jun 2019).
- [9] Bradford, D. F. (1975). “Constraints on Government Investment Opportunities and the Choice of Discount Rate”, *American Economic Review* 65(5): 887–99.
- [10] Broome, J. (1994). “Discounting the Future”, *Philosophy and Public Affairs* 23(2): 128–56.
- [11] Chetty, R. (2006). “A New Method of Estimating Risk Aversion”, *American Economic Review* 96(5): 1821–34.
- [12] 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).

- [13] 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).
- [14] Cotton-Barratt, O. (2015). “Allocating risk mitigation across time”, Future of Humanity Institute Technical Report #2015-2.
- [15] 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).
- [16] Council of Economic Advisers to President Barack Obama (2017). “Discounting for Public Policy: Theory and Recent Evidence on the Merits of Updating the Discount Rate”. URL: https://obamawhitehouse.archives.gov/sites/default/files/page/files/201701_cea_discounting_issue_brief.pdf (accessed 23 July 2019).
- [17] 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).
- [18] Drupp, M. A., M. C. Freeman, B. Groom, and F. Nesje (2018). “Discounting Disentangled”, *American Economic Journal: Economic Policy* 10(4): 109–34.
- [19] Gandelman, N., and R. Hernández-Murillo (2014). “Risk Aversion at the Country Level”, St. Louis Fed Working Paper 2014-005B.
- [20] Giglio, S., M. Maggiori, and J. Stroebe (2015). “Very Long-Run Discount Rates”, *The Quarterly Journal of Economics* 130(1): 1–53.
- [21] Greaves, H., and W. MacAskill (2019). “The Case for Strong Longtermism”, Working Paper. URL: https://globalprioritiesinstitute.org/wp-content/uploads/2019/Greaves_MacAskill_The_Case_for_Strong_Longtermism.pdf (accessed 7 Oct 2019).
- [22] 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).
- [23] Giving What We Can (2013). “Frequently Asked Questions”. URL: <https://www.givingwhatwecan.org/about-us/frequently-asked-questions> (accessed 29 May 2019).

- [24] Farhi, E., and I. Werning. (2005). “Inequality and Social Discounting”, NBER Working Paper 11408.
- [25] Farhi, E., and I. Werning. (2010). “Progressive Estate Taxation”, *Quarterly Journal of Economics* 125(2): 635–673.
- [26] GiveWell (2019). “2019 GiveWell Cost-Effectiveness Analysis – Version 6 (public)”, URL: https://docs.google.com/spreadsheets/d/1zLmPuddUmKsy3v55AfG_e1Quk-ngDdNzW-FDx0T-Y94/edit#gid=1034883018 (accessed 28 December 2019).
- [27] Giving USA (2019). “Giving USA 2019: The Annual Report on Philanthropy for the Year 2018”.
- [28] Gollier, C. (2016). “Evaluation of long-dated assets: The role of parameter uncertainty”, *Journal of Monetary Economics* 84: 66–83.
- [29] Gollier, C. and M. Weitzman (2010). “How should the distant future be discounted when discount rates are uncertain?”, *Economics Letters* 107(3): 350–3.
- [30] Gwern (2012). Untitled comment. URL: <https://www.lesswrong.com/posts/TYAA4iNCFaDvPD6gB/a-rational-altruist-punch-in-the-stomach#LdZ3Rt2Rp5iZ3TBYP> (accessed 28 December 2019).
- [31] Hanson, R. (2013). “If More Now, Less Later”, Overcoming Bias. URL: <https://www.overcomingbias.com/2013/04/more-now-means-less-later.html> (accessed 28 December 2019).
- [32] Hanson, R. (2018a). “Long Legacies and Fights in an Uncaring Universe”, Overcoming Bias. URL: <https://www.overcomingbias.com/2018/10/long-legacies-and-fights-in-an-uncaring-universe.html> (accessed 28 December 2019).
- [33] Hanson, R. (2018b). “Long Legacies and Fights in a Competitive Universe”, Overcoming Bias. URL: <https://www.overcomingbias.com/2018/11/long-legacies-and-fights-in-a-competitive-universe.html> (accessed 28 December 2019).
- [34] Hanson, R. (2018c). “Long Views Are Coming”, Overcoming Bias. URL: <https://www.overcomingbias.com/2018/11/long-views-are-coming.html> (accessed 28 December 2019).
- [35] Kaldor, N. (1957). “A Model of Economic Growth”, *The Economic Journal* 67(268): 591–624.
- [36] Klein, P. and J. V. Ríos-Rull (2003). “Time-consistent Optimal Fiscal Policy”, *International Economic Review* 44(4): 1217–45.
- [37] 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).

- [38] Koopmans, T. C. (1965). “On the Concept of Optimal Economic Growth”, in *The Econometric Approach to Development Planning* (Amsterdam: North-Holland Publishing Co.): 225–87.
- [39] Koopmans, T. C. (1967). “Objectives, Constraints, and Outcomes in Optimal Growth Models”, *Econometrica* 35 (1): 1–15.
- [40] MacAskill, W. (2015). *Doing Good Better* (New York: Random House).
- [41] Mankiw, N. G. (2015). “Yes, $r > g$. So what?”, *American Economic Review: Papers & Proceedings* 105(5): 43–7.
- [42] Matthews, D. (2018). “You have \$8 billion. You want to do as much good as possible. What do you do?”, Vox. URL: <https://www.vox.com/2015/4/24/8457895/givewell-open-philanthropy-charity> (accessed 21 Dec 2019).
- [43] Millner, A. and G. Heal (2018). “Discounting by Committee”, *Journal of Public Economics* 167: 91–104.
- [44] Moss et al. (2019). “EA Survey 2018 Series: Cause Selection”. URL: <https://forum.effectivealtruism.org/posts/hP6oEXurLrDXyEzcT/ea-survey-2018-series-cause-selections> (accessed 21 Dec 2019).
- [45] National Aeronautics and Space Administration (2007). “Near-Earth Object Survey and Deflection: Analysis of Alternatives”, Report to Congress. URL: https://www.nasa.gov/pdf/171331main_NEO_report_march07.pdf (accessed 27 Dec 2019).
- [46] National Philanthropic Trust (2019). “The 2019 DAF Report”. URL: <https://www.nptrust.org/reports/daf-report/> (accessed 21 Dec 2019).
- [47] Nordhaus, W. D. (2011). “The Economics of Tail Events with an Application to Climate Change”, *Review of Environmental Economics and Policy* 5(2): 240–57.
- [48] Olson, M., and M. Bailey (1981). “Positive Time Preference”, *Journal of Political Economy* 89(1): 1–25.
- [49] 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).
- [50] Parfit, D. (2011). *On What Matters, Vol. 2* (Oxford: Oxford University Press).
- [51] Piketty, T. (2014). *Capital in the Twenty-First Century*, translated by A. Goldhammer (Cambridge, MA: Harvard University Press).

- [52] Pindyck, R. S. (2002). “Optimal Timing Problems in Environmental Economics”, *Journal of Economic Dynamics & Control* 26: 1677–97.
- [53] Ramsey, F. (1928). “A Mathematical Theory of Saving”, *The Economic Journal* 38(152): 543–59.
- [54] Simon, L. K. and M. B. Stinchcombe (1989). “Extensive Form Games in Continuous Time: Pure Strategies”, *Econometrica* 57(5): 1171–214.
- [55] Tarsney, C. (2019). “The Epistemological Challenge to Longtermism”. Unpublished.
- [56] Tran, B. R. (2016). “Divest, Disregard, or Double Down?”, *Federal Reserve Board Finance and Economics Discussion Series* 2017-42.
- [57] Vivalt, E. (2019). “Give Later”. URL: <http://evavivalt.com/2019/07/give-later> (accessed 13 Aug 2019).
- [58] von Below, D. (2012). “Optimal Carbon Taxes with Social and Private Discounting”, Working Paper, SURED Conference, Ascona, Switzerland.
- [59] World Bank (2019). “World Bank national accounts data, and OECD National Accounts data files.” URL: <https://data.worldbank.org/indicator/ny.gdp.mktc.cd> (accessed 4 Nov 2019).