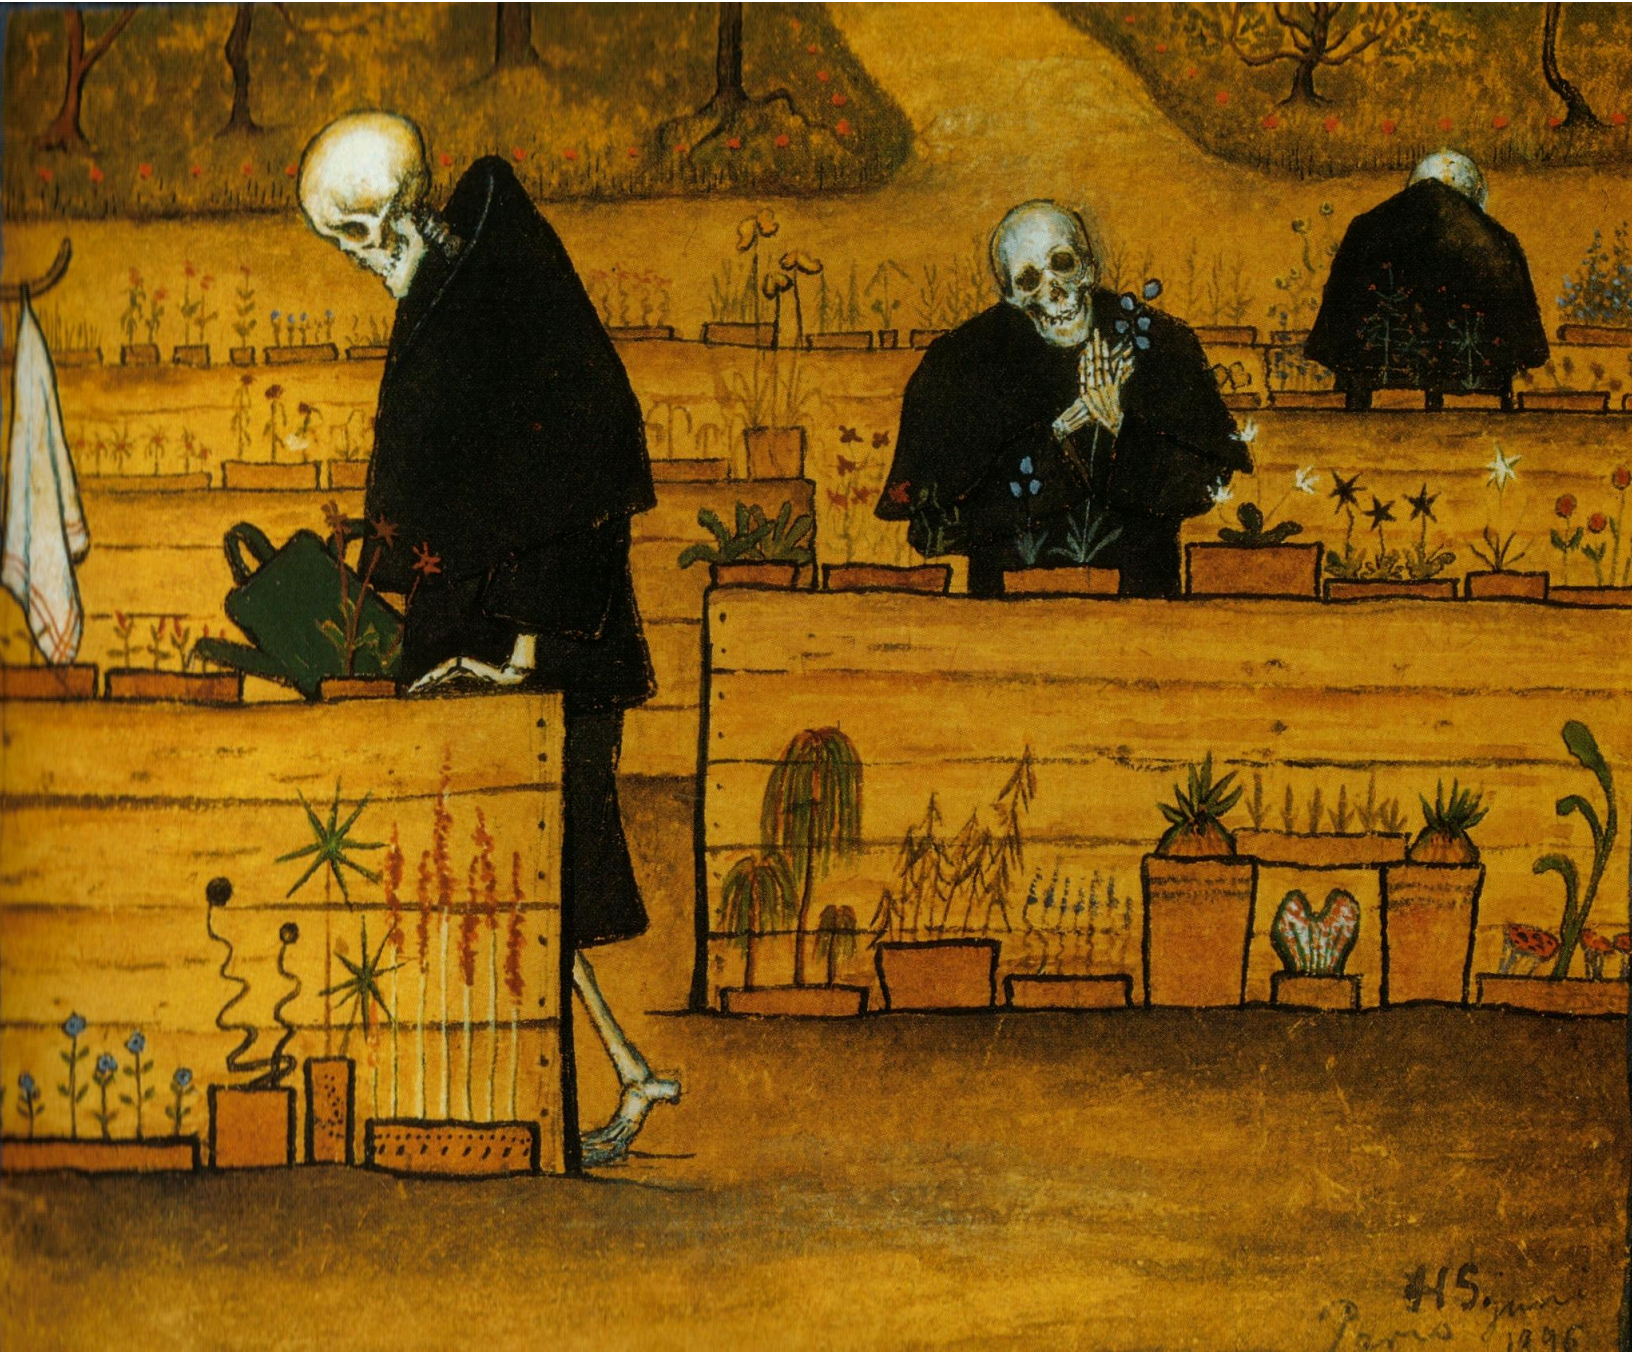

Patient Philanthropy in an Impatient World

Philip Trammell[1]

September 2021 (last updated: November 2022)

[1] Global Priorities Institute and Department of Economics, University of Oxford.
Contact: philip.trammell@economics.ox.ac.uk. I thank the Collison Family Office for funding the writing of this report and Adrian Ivashkiv and Luis Mota for assistance.

Contents

Summary

When people decide how to spend their money, and what fraction of it to invest for future spending, one consideration they take into account is their *pure time preference*. That is, they may value future benefits less in part simply because they are in the future. Another consideration is (lack of) *intergenerational altruism*. If someone dies before consuming his invested resources, the resources will instead be consumed by his heirs; and, in general, people value their heirs' consumption less than their own. We might call these two considerations "impatience".

People spend more and save less than they would if they had no pure time preference and were longer-lived. Also, their expenses generally consist of goods with more long-term costs and fewer long-term benefits: patient people spend less on sweets and video games and more on vegetables and education. Finally, the starkest difference between patient and impatient behavior should concern saving for future expenses with further-future benefits. A patient person is bound to be much more likely than an impatient person to save his dessert money so that he can later afford a healthier diet or better schooling for his children.

Firms, governments, and philanthropists likewise face the problems of how to spend on their respective beneficiaries today and how much to invest for their beneficiaries' heirs or future selves. Like individuals, they typically make these decisions impatiently.

By contrast, an impartial philanthropist may not be inclined to discount benefits to her beneficiaries purely because they are in the future, and she might care about present and future generations equally. That is, she may be patient. From her perspective, the world spends too much of its wealth each year and saves too little. Of what is spent, too much is spent on standard consumption goods and too little is spent on, say, basic research that will not bear fruit for a long time. Investment for the sake of future (say) basic research is most underfunded of all.

Therefore, **a patient philanthropist can typically do more good (from her perspective) by investing for the sake of future spending than by spending immediately.** She should only begin spending under two circumstances. First, she should spend once she, and any other patient funders in an area she wishes to support, have grown wealthy enough relative to the area's impatient funders that even the impatient are spending at *less than the patient-optimal rate* for the collective (patient plus impatient) budget. Second, she should spend if she finds a fleeting opportunity to achieve sufficiently outsized long-term impact: if the vegetable shop, as it were, is having a steep enough sale.

This framework simplifies the patient philanthropist's problem. If she were trying to determine her ideal spending rate each year *in isolation*, she would have to base her decision on long-term forecasts of the interest rate, the economic growth rate, and many other variables. But as a small patient player in a large impatient world, she can often infer from impatient parties' spending behavior that her own optimal spending rate is zero.

By following the policy outlined above, the patient philanthropist can plausibly do **hundreds of times more good directly** than she would by spending her endowment on a

more typical, "impatient" schedule. This is true even though this policy may require investing for centuries before spending any significant sum, and even after accounting for the fact that doing so runs a high risk that the invested resources are expropriated, or spent in ways the founding philanthropist would not endorse. Furthermore, by lowering interest rates or raising growth rates (however slightly), **long-term investment can yield significant indirect benefits.**

A patient philanthropist should invest her resources somewhat differently from an impatient or self-interested party. Because of her patience, she should be willing disproportionately to **invest in risky and illiquid assets.** Because of her philanthropy, she should be willing to make investments with below-market private returns if they produce large enough social benefits, and she should try to choose assets that pay off most in the event that her philanthropy is most needed. Though the latter two investment considerations apply to some extent to all philanthropic investors, however, they apply more weakly when the philanthropic investor is more patient. When a significant part of a philanthropist's assets are to be invested until far in the future, it is most important simply to maximize private investment returns.

Finally, a patient philanthropist should consider "investing" in **growing the community of people who are passionate about her own philanthropic objectives and equipped to execute on them.** In part this is because such individuals can invest and spend their own capital as she would, effectively increasing the size of her endowment. Indeed, recruiting individuals to use their wealth as one would oneself is better modeled as a kind of financial investment than as philanthropic spending. However, a more significant benefit to growing the community of actors sympathetic to one's cause, alongside one's endowment, is that labor drawn from this community is complementary to one's philanthropic capital (especially once this capital has passed to future generations). For instance, the larger one's endowment, the more valuable it is to have a large pool of competent and ideologically sympathetic grantmakers.

1. Introducing the scheduling problem

If we have capital to spare, we might ask how we can best use some of it to have a positive impact on the world. We might consider both direct spending to improve the world ("giving") and investing in profitable projects that one believes will have positive impacts on the world ("impact investing"). Following Roth (2021), let us refer to this pair of projects as "philanthropy".

Thoughtful, strategic philanthropy is a complex project, raising many questions with no clear answers. However we engage with the project, however, a central question we must confront is the question of disbursement scheduling. Indeed, once we have resolved to do good, we face this question immediately. Implicitly or explicitly, we must decide at every moment how much of our wealth, if any, to spend now and how much to invest for some kind of impact in the future. We may intend never to disburse, if we judge that our investments themselves have "social returns" high enough that the private returns are best wholly reinvested. That, too, would be a scheduling decision.

An intimidating array of considerations bears on the scheduling question. Investing funds for future use gives us time to learn more about how to do good, including by reading others' ongoing research and learning from their experiences; it retains us the freedom to spend or invest in future funding opportunities, which may be more valuable than any available today; and it avails us of the miracle of compound interest. On the other hand, it shifts our funds toward efforts that will, if economic growth continues, likely be better funded than comparable efforts today; it risks leaving investment and disbursement decisions in the hands of people who do not share our values, whether these are strangers in future generations or simply our future selves; and it runs the risk of delaying some or all our potential impact until our funds are expropriated or our civilization collapses. Spending now teaches us how to identify and implement worthwhile projects in the future, through trial and error; if we fund sufficiently relevant research, it gives us direct insight along the same lines; and it triggers a chain of flow-through effects that may themselves compound over time. The list of considerations goes on.

1.1 *Avoiding the issue*

Despite (or because of) the complexity of the problem, three simple disbursement schedules are especially common among large philanthropists. First, ("pure") impact investors, by definition, never disburse. Second, large givers often aim to spend everything during their lives, or within a short period after their deaths—with spending typically skewed toward the end of the period in question. Finally, it is also common for large givers to establish a perpetual foundation that disburses roughly 5% of its assets per year, the legal minimum for American tax-exempt foundations.

Through most of the 20th century in the United States, perpetuities such as the Rockefeller, Carnegie, and Ford Foundations were the standard vehicles for disbursing large philanthropic

endowments and spent at roughly the legal minimum rate. In recent years, by contrast, it has become more common to practice and advocate "giving while living".[2] Calls for faster philanthropic spending have gained new prominence, recently culminating in the Accelerating Charitable Efforts (ACE) Act. Proposed in the US Senate, the ACE Act would impose disbursement requirements on donor-advised funds (DAFs)—tax-exempt, foundation-like investment accounts for smaller donors—and tweak the tax code to encourage faster spending by foundations. The bill was introduced by Senators Grassley (R-IA) and King (I-ME), but it was largely written by the Initiative to Accelerate Charitable Giving (IACG), a private consortium established in 2020.

Despite their wide and growing prevalence, calls for fast disbursement, or for any given disbursement scheduling policy or norm, have come with surprisingly little deep public analysis. Some people categorically oppose the directed funding of charitable efforts after the funder has died on the grounds that this arrangement is undemocratic.[3] For these people, perhaps, the question needs no further analysis. Even among those who do not accept this sort of blanket moral prohibition, however, standard practice appears to be to identify a single relevant consideration and then to proceed as if the entire question is settled. On one side of the debate, the IACG justifies its mission on the grounds that now is a time of "unprecedented needs", without discussing whether, for instance, investment returns might *outweigh* the extent to which the future is less needy—or, strictly speaking, even whether to expect that the future will be less needy.[4] On the other side, the Council on Foundations, a philanthropic interest group that largely lobbies against regulations on foundations, even more shallowly argues that stricter disbursement requirements "will not achieve … greater support for nonprofits and communities, now and in the future".[5] Cowen (2006, pp.36–37) offhandedly (though tentatively) argues that foundations should be required to spend down quickly, on the assertion that young foundations are more dynamic and effective than old ones. Single-issue cases for giving quickly can likewise be found throughout the statements made by Giving Pledge signatories. None of these parties appear to provide substantially more detailed justifications elsewhere for their respective positions on giving schedules.[6]

---

[2] Holcombe (2000) includes a particularly thorough documentation of the beginning of this transition. More recent commentaries on philanthropy, such as Reich (2018) and Callahan (2018), also touch on the trend, especially as it has manifested itself over the past few decades. The term "giving while living" was coined by Atlantic Philanthropies founder Chuck Feeney, and the practice is defended at length but highly informally in his foundation's 2017 publication by the same name. For a particularly stark example of advocacy for faster philanthropic spending in recent popular media, see Schleifer (2019).

[3] One Reverend John Haynes Holmes famously, in 1915, called long-lived foundations "repugnant to the whole idea of a democratic society".

[4] See Initiative to Accelerate Charitable Giving (2021), for instance.

[5] Council on Foundations (2021)

[6] None of this is to say that little is written on the question, only that it appears generally to be informal or one-dimensional. Estimates of the near-term disbursement increases that would be produced by a given disbursement minimum are particularly common, for instance. For representative samples from the existing discussion, see Leat (2016), Chapter 7; Pharoah and Harrow (2010); or Association of Charitable Foundations (2016).

The lack of formal analysis of the disbursement scheduling question among philanthropists and philanthropy commentators may be due in part to the question's neglectedness among academics. To my knowledge, Andreoni (2018) offers the only academic analysis to date of the welfare implications of disbursement requirements, and Trammell (2021) offers the only academic analysis of the implications for disbursement scheduling of "time preference heterogeneity", i.e. the fact that some funders care more about the future than others. (These will be discussed in Sections 5.3 and 3, respectively.) Much more work on these and related questions remains to be done.

In short, the disbursement scheduling question may both deserve and be amenable to more careful scrutiny than it currently receives.

1.2 *Adding the rates*

On attempting to grapple seriously with the question of whether to begin giving immediately, it is common to begin by listing the trends at play: the interest rate, the rate at which economic growth will decrease the marginal value of philanthropic spending, and so on. Each of these trends corresponds with a rate that is positive or negative, depending on whether it increases or decreases the impact of spending in the future relative to the present. The rates are estimated and summed. Then, in the words of Wayne Chang, a professor of finance at Minerva University, if the "sum is above [zero], it makes sense to give later. If it falls below [zero], it makes sense to give now." Let us call this approach "adding the rates", or AR.[7]

This approach is not correct. In particular, it fails for three reasons: because the estimated rates could change in the future, because there can be diminishing returns to giving at any one time, and because the optimal spending plan of any one philanthropist both depends on and affects the spending of others. The next three subsections will explore these problems further. Nevertheless, AR is a start; and the wideness of its appeal, even among the financially literate, suggests that it is worth beginning here and working toward a more robust framework by exploring AR's faults and correcting them.

1.3 *Timing the market*

The most straightforward issue with AR is that it fails to account for the possibility that the rates in question will change so as to reverse the value of waiting relative to that of disbursing. This failure amounts to ignoring the *option value* that is kept only by waiting to disburse. That is, when one spends resources, one locks in the timing of their expenditure; but when one waits, one retains the option either to spend in the next period or to spend in any subsequent period.

[7] Chang expresses this conclusion in a 2020 forum comment. Other attempted implementations of AR from the Effective Altruism community, beyond the post to which this comment replies, include Hurford (2013), Robinson (2016), and Dickens (2019). I do not mean to suggest that these relatively informal analyses represent their authors' best judgment, only to demonstrate the commonness of the corresponding intuition, at least among (perhaps small) donors beginning to think through the scheduling problem thoroughly.

The risk is asymmetric. When AR produces a positive number—i.e., produces a recommendation not to disburse—this does not lead us astray. When it produces a negative number, however, disbursing may still be a big mistake.

To see this, suppose for illustration that the only considerations at play were

1. the interest rate (making it better to give later) and
2. the rate at which economic growth is decreasing the marginal value of philanthropic spending (making it better to give sooner).[8]

If (1) currently exceeds (2), marginal funds are better spent in the near future than spent immediately. Even if we have not determined whether it is best to spend in the near future or the further future, we know that it is not optimal to spend in the present.

If (2) currently exceeds (1), on the other hand, we know only that marginal funds are better spent now than in the *near* future. Suppose, for example, that (2) will exceed (1) only for five years, and that (1) will then exceed (2), to a similar extent, for ten years. In this case, disbursing now does more good than disbursing in one year, better still than disbursing in two, and so on until year five. But then the trend reverses: it is slightly less bad to disburse in year six than in year five, and by year ten or so, the gift will be doing about as much good as would have been accomplished by giving immediately. By year fifteen, waiting will have paid off, perhaps substantially.

This is more than a theoretical possibility. As discussed in Section 2.6, there is reason to believe that even if (2) currently exceeds (1), (1) will exceed (2) in the long run.

In short, adding the rates is not sufficient. Our decision about whether to spend now is sensitive not only to the values of the relevant rates as we observe them today, but also to our guesses about how they will evolve over the entire future. Adjusting AR for the possibility that rates can fluctuate gives us a new approach, which we might call "timing the market", or TM. Timing a philanthropic market is not like timing a typical financial market, because gifts, unlike financial positions, are generally irreversible. It is not enough to know whether a giving opportunity is more or less valuable than it will be tomorrow. We have one, irrevocable time to spend each (present-valued) dollar, so we have to do our best to spend each dollar when it is best to spend it, out of all the times the future has waiting.

1.4 *Balancing the portfolio*

TM does not solve all of AR's weaknesses. A remaining weakness shared by TM is the implicit assumption that each dollar spent at a given time has equal impact.

This assumption is reasonable for small funders. All else equal, spending by any one small donor is unlikely to drive down funding needs in some area by enough that the donor's final dollar spent has less impact than his first dollar. However, sufficiently large donors do face diminishing returns to spending at a given time. Therefore, if a large donor spent all her

[8] Note that this rate is not the growth rate itself, and may be higher or lower. For illustration, if the cost of doing a unit of good did not fall at all as people grew wealthier, then (2) would always equal zero.

resources at any single time, even if it were the best time to spend her *first* dollar, her final dollars spent at that time may not be put to good use.

That is, after accounting for diminishing returns, the disbursement scheduling problem really is one of scheduling, rather than timing. The goal is no longer to find "the best time" to spend. Instead, it is to find the set of times across which to spread one's budget, such that spending more at any one time is no better than spending more at any other, and spending at any other time is worse. The resulting optimal schedule will be sensitive to guesses about both present and future interest rates and so on, as under TM, but now it will also be sensitive to estimates of the rate at which impact diminishes as the spending rate at each time increases.

Given diminishing returns, one's optimal disbursement schedule can typically be found by determining the optimal *global* spending schedule (at least on the kinds of projects one wishes to fund), comparing it to the global spending schedule that one expects to observe in the absence of one's own contributions, and spending so as to shift the global spending schedule as close as possible to the optimum. We might therefore call this approach "balancing the portfolio", or BP.

As this framing emphasizes, when the returns to spending are not constant, one's disbursement scheduling problem cannot be solved in isolation. Just as the time at which someone should spend a given dollar depends on when he plans to spend the rest of his budget, the schedule on which someone should spend his budget depends on the schedules on which others spend theirs.

1.5 *Playing the game*

The last issue with the approaches above, including BP, is that they implicitly assume that the variables governing our spending problem are exogenous. In fact, they may depend on what we do.

For most relevant variables, it is reasonable to assume that how they change over time depends only negligibly on our own decisions, at least unless our wealth or spending has reached monumental proportions. Interest rates may rise and fall, for instance, but unless one is directing a substantial fraction of global capital, they will not change much in response to one's own disbursement rate or investment behavior. Still, dependencies of this kind may sometimes be important, especially when considering whether to attempt investments over sufficiently long timescales. For instance, when evaluating a (presumably intergenerational) plan to invest until one's fund has reached a value of ten trillion dollars before disbursing, it may be necessary to treat each year's expropriation risk not as exogenous but as increasing in fund size. This will be discussed further in Section 5.3.

One important variable, however, is likely highly sensitive to our own disbursement decisions, however small we are as funders. Indeed, in a certain sense, it is likely to be more sensitive to our decisions the *smaller* we are. This variable is the spending schedule collectively adopted by other funders.

As noted in the discussion of BP, the schedule on which you should spend depends on the schedules on which others spend. But the converse is also true: others' optimal spending

schedules, given their beliefs and goals, depend on your own spending schedule. Your spending can thus change others' behavior, producing an allocation of resources across time (and across projects at any given time) very different from the allocation you would expect before taking these game-theoretic issues into account.

If other funders are not strategic, in the sense that they do not adjust their spending schedules in light of one's own behavior, one can ignore this complication and simply, more or less, implement BP. But game-theoretic complications arise as long as there is *any* other player who is considering funding the same projects at the same time periods and who is taking their own disbursement scheduling problem seriously.

For illustration, suppose there is a single project to fund and two years in which to fund it. You have $10 to spend across the periods; another strategic funder—let's call her Gillian Bates—has $20; and non-strategic funders have $40. Before making your spending plans, the other parties are planning to allocate their resources like this:

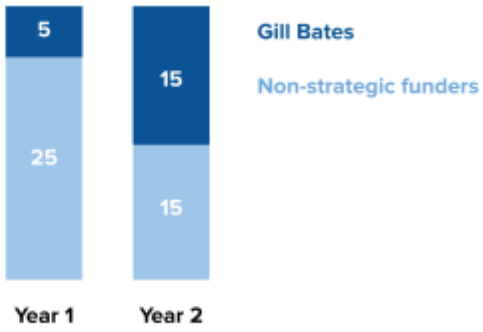

Figure 1

Non-strategic funders spend $25 in the first year and invest $15 for spending in the second, regardless of others' behavior. Gill believes that allocating funds for the first year is just as important as allocating them for the second, keeps track of the funding coming from others, and allocates her $20 budget accordingly. From her perspective, she can balance the collective portfolio.

You believe that funds are better spent in year 2. If you could, you would like to use your resources to raise the total budget allocated to year 2 from $30 to $40. On seeing that you have no plans to disburse in year 1, however, Gill can adjust her own schedule so as to keep the resources allocated to each period equal.

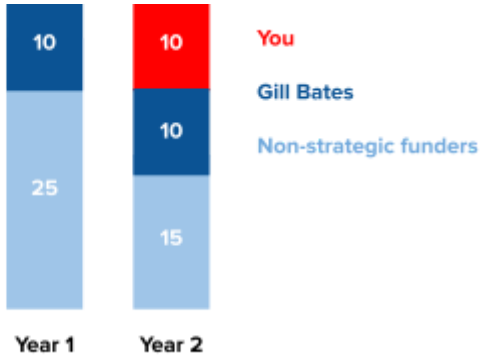

Figure 2

Indeed nothing you do in this context can have any impact on the collective spending schedule. If you had more than $30 to spend, you would be able to skew collective spending toward year 2, since Gill cannot cut second-year spending by more than $15. Likewise, if you had more than $10, you could skew collective spending toward year 1, if you chose. But if you have $10 or less to spend on this project, you ultimately only have one scheduling option: to allocate $5 to each year.[9]

The issue of philanthropic displacement, or "crowding out", is not unique to the problem of scheduling. Simply replace "years 1 and 2" with "projects 1 and 2", and the reasoning of the example above still applies. Indeed, crowd-out issues are much more extensively studied, and are often easier to infer, in the single-period/multi-project context. In October 2017, for example, Bill Gates announced his foundation's plan to increase funding for educational initiatives in the US: having spent $1 billion in the area over the previous 17 years, they would now spend almost $1.7 billion over the next five.[10] This announcement came on the heels of an August 2017 post in which Gates touted the "remarkable", "miraculous" progress toward another of his foundation's goals—malaria eradication—made over the preceding decade, which he attributed largely to increases in global funding for the effort.[11] A natural conclusion is that hundreds of millions of dollars that outside funders had intended to direct to the prevention of malaria deaths were actually directed, in the final accounting, to the aid of American students. In short, even when it is not directly observed, the phenomenon of crowd-out complicates any project in which a philanthropist might hope to engage, so long as he and the other potential funders of the project in question do not have identical preferences.

Nevertheless, as I will argue in Section 2, time preference differences are a particularly important and pervasive class of preference differences philanthropists face. They also present complexities, and opportunities for fruitful collaboration among funders, that do not appear in a single-period setting. We will explore some of these complexities and opportunities in a highly simplified setting in Section 3, in the context of "patient giving"—strategic disbursement

[9] For evidence on intertemporal crowding out in the context of contributions to university endowments, see Oster (2001).
[10] Gates (2017b)
[11] Gates (2017a)

behavior by patient actors in the presence of impatient co-funders—and in Section 4, in the analogous context of patient impact investing. Section 5 will conclude by discussing the lessons of history for the challenges and prospects of implementing patient philanthropy of either kind.

2. Time preference

The scheduling problem we face, when we aim to do the most good with our money, cannot possibly be solved. It cannot even be written down. It is less like an arithmetic problem than like a monstrous game of chess, played by millions of players, across countless generations, on a vast grid that is always changing. And because we only play it once, there are some things we cannot learn entirely by experience. If we are to approach a better understanding of how to play this game well, we have to start by studying the dynamics produced by just a few of its rules, in a narrow range of simple situations.

If everyone in the world had the same beliefs and goals, attempts to achieve impact would not require "playing a game". We would still collectively face a difficult optimization problem in allocating the world's resources across projects and across time; but our only concern would be the construction of that "portfolio", not the reactions of other parties to our own spending decisions. To identify the most important game-theoretic concerns, therefore, we must begin by identifying the ways in which our preferences most differ from those of other parties. And one of the starkest and most consequential ways in which our preferences may be unusual, if we seek to do the most good, is that we may not be inclined to discount future wellbeing.

2.1 *Individual time preference*

We all make countless decisions that require us to weigh our own wellbeing at one time against our wellbeing at another time. We must decide how hard to work in school in exchange for a better job, or how closely to watch our diets in exchange for better health. Most straightforwardly, and most relevantly for our purposes, we must decide how much of our wealth to consume each year and how much of it to invest.

Different people make these tradeoffs differently, and people generally make them differently in different contexts.[12] At least when allocating spending across time, however, people seem to discount their own future wellbeing by at least 2–3% per year on average.[13] That is, in situations in which individuals know that they will consume no more or less next year than this year (so a little more money might reasonably be expected to "buy just as much wellbeing" for

[12] Evidence for this context-dependence is summarized and emphasized by Frederick et al. (2002).

[13] See for instance Alan and Browning (2010). They also offer a helpful (though dated) review of the pre-existing econometric and experimental literature on the distribution of time preferences across individuals and households, on pp. 1252–3.

them next year as this year), the market-clearing risk-free real[14] interest rate is not zero, but a few percent higher.[15]

In more realistic settings, of course, many other factors determine interest rates as well. When people expect their consumption to increase over time, for example—the typical case, given positive economic growth—the interest rate must be higher, because people do not have as much need for wealth in the future. But regardless of the other factors at play, to leave people indifferent between consuming more now and investing to consume more next year, they must be compensated with a few percent in interest *just for the wait*.

This is not surprising to anyone familiar with the human experience. Imminent rewards and punishments weigh on us more heavily than distant ones. To some extent, this time preference can be attributed to mortality risk. However, since the average investor faces substantially less than a 2–3% chance of death per year, and since people typically care at least somewhat about the size of the bequests they leave their heirs, some of the time preference we observe in people's investment behavior seems to constitute "pure" time preference.[16]

If people discount their (or their heirs') future wellbeing at a constant rate, whatever that rate is, they discount *exponentially*. Often, however, the rate at which people discount each year appears to fall the further the year is in the future. They might currently be indifferent between 1 unit of wellbeing this year and 1.03 units next year, for instance (exhibiting a discount rate of 3%), but between 1 unit in ten years and 1.01 units in eleven years (exhibiting a discount rate of 1%). People with preferences like this discount *non-exponentially*. Hyperbolic discounting, for instance, is a particularly common, and now particularly well-known, pattern of non-exponential discounting.

Someone with non-exponential time preferences is "time inconsistent", in that she will continually make decisions that she would not have endorsed in the past. She may wish today to trade a reward in 10 years for a 2% larger reward in 11 years, but once year 10 arrives, she will evaluate this tradeoff less patiently. People with exponential time preferences do not have this issue. For this reason, economists typically consider exponential discounting "rational" and other patterns of time preference "irrational".

The feature of individual time preference on which I wish to focus, however, is not irrationality but simply impatience, whatever form it may take. Furthermore, if people cannot tie their hands to spend more patiently in the future, as they very often cannot, they will always spend in light of the discount rate they use to compare benefits across periods in the immediate future (i.e., in the example above, 3% per year). Interest rates will be set accordingly. For simplicity, therefore, I will assume throughout most of the discussion below that people discount at a constant rate.

---

[14] That is, after inflation. Throughout this document, we will refer to all variables in real terms.

[15] Note that this is only the risk-free interest rate that will obtain in a "risk-free world": a situation in which, as stipulated, individuals know that their future consumption will not rise or fall for other reasons. Given background risk, risk-free investments may command lower or even negative interest rates, because they are serving to some extent as insurance.

[16] It has been proposed that this inclination to discount "excessively" may be due to the higher mortality rates we faced in the distant past, especially outside old age.

2.2 *Firm time preference*

Firms also have time preferences. Public firms (or their managers) are often accused of being too "short-termist", pursuing short-term profits at the expense of longer-term viability. Numerous laws, regulations, and standard business practices have been blamed for this short-termism. It is argued, for instance, that quarterly earnings requirements incentivize managers to avoid ventures that pose short-term risks to the bottom line, even when these can be expected ultimately to maximize shareholder value.[17]

To some extent, these attributions may be warranted. A subtler but much more fundamental driver of (at least public) firm short-termism, however, is simply the phenomenon of the section above: individual impatience and its effect on interest rates.

To illustrate this, suppose a CEO must choose between two risk-free $100 capital investments: one offering a return of 5% after one year and one offering a return of 10% after three years. These returns will be paid as dividends to the firm's shareholders as soon as they are produced. Finally, suppose the risk-free interest rate is 4%, of which 3% is due to individuals' time preference.

As the 4% interest rate reveals, individuals are generally indifferent between about $101 immediately and $105 next year. That is, the first investment will turn $100 into a future dividend that shareholders value at about $101. It thus raises the value of the firm by about $1. By contrast, because $98 \times 1.04^3 \approx 110$, individuals are generally indifferent between $98 immediately and $110 in three years. The second investment would thus lower the value of the firm by about $2.

Given a prevailing interest rate of 4%, this is entirely reasonable. Even shareholders who would prefer to consume an extra $110 in three years than to consume an extra $105 in one year will prefer the one-year dividend of $105, because they can invest this sum elsewhere for the following two years at 4% interest, producing $114 rather than $110 in year 3.

In a more patient world, by contrast, people would invest more on the whole, and this would lower the interest rate. As stipulated, a fully patient world here would invest until the interest rate had fallen all the way to 1%. Faced with the above pair of internal investment opportunities, then, our CEO would find that individuals are generally indifferent between about $104 immediately and $105 next year, so that the first investment raises the value of his firm by about $4. But because $107 \times 1.01^3 \approx 110$, the second investment would do even better: it would raise the value of the firm by about $7.

In short, what it *means* to maximize shareholder value is ultimately defined by the time preferences of the society in which a firm is embedded. In an impatient world, even when managers are flawlessly executing their fiduciary duties, the impatience reflected in prevailing interest rates drives them to allocate their resources impatiently.

2.3 *Government time preference*

Like firms, governments must regularly allocate capital among projects with different streams of costs and benefits. When these allocation decisions are made directly by referendum, or by a vote of elected representatives, they presumably tend to reflect individuals' preferences,

[17] See e.g. Koller et al. (2017) and Dimon and Buffett (2018).

including their time preferences. Also, of course, as in the case of firms, institutional variables like election frequency can encourage or discourage long-term thinking to some extent, and are often blamed for doing the latter.

However, most public resource allocation decisions in most large developed-world governments are not set by law or direct democracy but by bureaucratic administration. These decisions even more explicitly incorporate individuals' time preferences, as reflected, again, in interest rates.

The standard administrative procedure for evaluating whether a public project is worth funding is known as cost-benefit analysis. It has two steps. The first is to convert all its costs and benefits into monetary figures, by estimating how much the affected parties would be willing to pay for the benefit, or to avoid the cost, when it arises. The second is to determine whether the project's *net present value*, i.e. the sum of the present values of all the costs and benefits, is positive. The present value of a monetary cost or benefit accruing at some future time is simply the amount today that would have to be invested to equal the given sum at the given time.[18]

Consider, for example, a municipality's decision of whether to build a bridge. The bridge will cost $100 million today. It will collect no tolls, but its users' willingness to pay for it—the maximum that *could* have been extracted in tolls—is projected to equal $5 million per year for each of the next 40 years, after which the bridge will not be usable.

Suppose the relevant interest rate is 4%. The present value of the first year of usage is then about $4.8 million, since $4.8 \times 1.04 \approx 5$. The present value of the second year of usage is about $4.6 million, since $4.6 \times 1.04^2 \approx 5$. Summing all four decades of benefits, their present value is $99 million. The bridge is thus judged not to be worth taxing $100 million to build, even though in absolute terms it would produce twice as much "wealth" ($5m × 40 = $200m) as it would cost.

As in the firm context, *given* an interest rate of 4%, this is the right decision. Taxing citizens to build the bridge would amount to redirecting $100 million from higher-return private projects to a lower-return public project, so it would do more harm than good even from a patient perspective. In a more patient world, however—one with a high enough saving rate that lower-return private projects *were* funded, and the interest rate were only, say, 1%—the bridge would be judged well worthwhile. Again, what it means for a public project to be of "positive net present value", or "pass cost-benefit analysis", is governed by the prevailing interest rate, and so ultimately governed by the patience of the average private investor saving for his own future.

How governments typically evaluate public projects is thus closely analogous to how firms evaluate private projects, as described in the previous subsection.

An interesting implication of this reasoning is that, if a policymaker has a lower rate of time preference for her constituents' (and their heirs') future welfare than the constituents have themselves, she should tax her constituents' consumption, and subsidize their investment, until they are acting as patiently as she would like (Barrage, 2018). That is, she should subsidize investment until interest rates have fallen to the rates that would obtain in a patient world. This would single-handedly induce individuals, firms, and cost-benefit analysts in government to

[18] See for example the policy governing discounting in cost-benefit analysis by the US federal government, given by US Office of Management and Budget (2019), or that employed by the UK government, given by HM Treasury (2020).

allocate capital in the patient-optimal way. Of course, however, such a policy would presumably not be sustainable in a democratic society.

Many projects promise to save or end lives in the future as well. Spending to prevent the spread of a deadly contagious disease, for instance, might predictably save lives for some time after the spending. Likewise, spending on efforts to mitigate the risks of climate change promises to save lives that have not even yet been born.

These benefits, too, are priced. The monetary value of a project that will, in a given year, save one unknown life out of (say) 300 million Americans is determined to be the average citizen's willingness to pay to lower their mortality risk that year by 1 in 300 million, multiplied by the 300 million individuals who will enjoy this risk reduction. This is known as the *value of statistical life*, or VSL. The monetary value of a project that will save multiple lives is simply the VSL multiplied by the number of lives expected to be saved.

Then, in an act of grim consistency, these benefits too are discounted. If the interest rate is 4%, and is projected to remain around 4%, a dollar-valued mortality risk reduction counts for 4% less for each year the life-saving is delayed. Some of this—1%, under the stylized numbers we have been using—is for reasons even a fully patient actor might endorse. As people get richer, for instance, they will be willing to pay more to reduce their mortality risk because they value money less; this should not inflate *our* willingness to pay to save their lives. But the remainder is time preference. The casual impatience of today's secretaries and dentists and mechanics, when deciding how much to save for their own retirement, putatively lets their representatives compare the disvalue of a child's death in 150 years with that of a child's death in 75 years and conclude that, compared to the latter, the former is almost meaningless.[19]

Similar logic is used to justify meager government spending on efforts to avert existential catastrophe. As Nordhaus (2009) observes, the US Congress at the time appropriated about $4m per year to NASA to track asteroids posing risks to human civilization, so that if one were found on a worrisome trajectory, we would have more time to prepare. More has been done since,[20] but more could be done yet: by spending $1b per year on better monitoring (including of comets, which are more difficult to track) and on deflection technology, it is estimated that we could reduce the chance of a disastrous collision by at least 90%. The B612 Foundation and the

[19] A representative illustration of this approach to cost-benefit analysis with respect to mortality risk can be found in the US Department of Transportation's Highway Safety Benefit-Cost Analysis Guide. An even more explicit example of intrinsically discounting far-future deaths is DICE-EMR (Bressler, 2021), the first integrated assessment model of climate change to account explicitly for mortality risk. It returns policy recommendations on the basis of discounting the welfare losses of future deaths at 1.5% per year. Note that this is in addition to the flow-through benefits of earlier lives for economic growth in future periods, which standard administrative cost-benefit analyses ignore but which are (at least intended to be) captured in DICE-EMR. Finally, until the early 2000s, reports on the Global Burden of Disease—a widely-used perpetual study, maintained by the Institute for Health Metrics and Evaluation (IHME), on the burdens imposed by all major causes of death and disability worldwide—discounted a disease's future costs in "quality-adjusted life-years" (QALYs) at 3% per year. The IHME has since switched to simply reporting the schedule of forecast QALY costs, so that the user can aggregate these cost-streams according to her own rate of time preference, but future QALY costs are still usually discounted at some positive rate for policy purposes. See e.g. Chen et al. (2015), Section 4 for a brief overview of the history, and Smith (2017) for a more extended account.

[20] Asteroid monitoring expenditures have increased by about an order of magnitude over the past decade; see NASA's Planetary Defense Coordination Office for details.

Planetary Society are nonprofits aiming to fill this gap, both of which could use more funding.[21] On Nordhaus's rough calculation, however, after valuing the benefits of the future of human civilization in consumption units and discounting them by the interest rate, substantial further increases in near-Earth object defense outlays would not pass cost-benefit analysis.[22]

2.4 *Philanthropist time preference*

For our purposes, what distinguishes philanthropists and governments from firms and individuals providing for themselves is that philanthropists and governments allocate capital so as to help other parties who are already spending on themselves—or, more generally, so as to fund causes which others are already to some extent funding. What distinguishes philanthropists from governments, in turn, is that a philanthropist makes his allocation decisions in light of his own moral beliefs or preferences, not those of any constituents.

There is a sense in which firms' investment decisions are undertaken on behalf of other parties too, of course. Firms, however, are only empowered to act so as to satisfy the preferences of their owners. As noted in Section 2.2, this will generally entail maximizing the firm's market value, and in an impatient world this in turn will entail acting with a certain degree of impatience (regardless of the owners' own time preferences). Likewise, a policymaker acts on her constituents' behalf, but in a democratic society she is charged with setting policy that satisfies her constituents' preferences. As we have seen, the behavior of firms and governments is thus largely determined by the rate of time preference that the average individual exhibits when making investment decisions for herself and her family. Philanthropists are essentially the only parties able to allocate substantial resources for socially valuable purposes on the basis of idiosyncratic attitudes toward time.

This is not to say that all philanthropists use this flexibility wisely. Holcombe (2000, p.237) asserts that many philanthropists spend too slowly than would be ideal, to accrue the prestige that comes with managing a large endowment. Robin Hanson, on the other hand, maintains (as reported by Alexander (2013)) that many philanthropists spend too quickly, to reap the social approval that comes with visible acts of charity. There is likely truth to both claims. At least some philanthropists, however, are surely willing to think carefully and sincerely about how to use their freedom with respect to time in the most morally appropriate way. These are the primary intended audience of this report.

Many argue that valuing future wellbeing no less than present wellbeing is a central part of what it means to seek to "do the most good".[23] It does not seem easy to justify a claim that the pain of a pang of hunger in 2025 detracts more from the moral value of the universe, as it were, than the

[21] For estimates of the cost-effectiveness of funding these organizations, see Newberry (2021).

[22] Some have argued that deflection technology would also inevitably allow omnicidal (or genocidal) actors to steer approaching bodies *toward* Earth (or particular regions), and that this risk outweighs the reduced risk of natural catastrophe (see e.g. Sagan and Ostro, 1994). Since this claim is disputed, however (Ord, 2020, ch.3), and since it does not appear to be Nordhaus's concern, Nordhaus (2009) serves to illustrate the relevance of time preference to federal spending decisions on existential risk reduction efforts. Ord also points out that the slow ramp-up of spending on monitoring alone has already exposed us to some avoidable existential risk, though this foregone reduction has been less than one in one million.

[23] See e.g. Cowen and Parfit (1992) and Cowen (1992).

pain of an equal hunger pang in 2125. It seems even more absurd to claim that it detracts less than the pain of an equal hunger pang in the past, as discounting for time would imply.

Furthermore, regardless of whether we accept this as a moral argument, many are simply inclined to be more patient—perhaps even fully patient—when making welfare tradeoffs for others than they would be when making such tradeoffs for themselves,[24] or (especially) when weighing benefits to themselves against benefits to their distant descendants.

For these reasons or for any other, a philanthropist may find himself a patient actor in an impatient world. More formally, consider the stylized figures used above: an interest rate of 4%, of which 3% is attributable to individual time preference. Though individuals are indifferent between an additional $1 of consumption this year and $1.04 for their or their heirs' consumption next year, a fully patient philanthropist in this setting is indifferent between augmenting others' consumption by $1 this year and by $1.01 next year. The interest rate in an impatient world thus strikes the fully patient philanthropist as favorably high.

Alternatively, instead of being "fully patient", a philanthropist may wish to respect her beneficiaries' *pure* time preferences. That is, she may feel that, if someone puts less weight on his own wellbeing the further it is in the future, it would be undesirably paternalistic for her not to do the same. Even so, she may not wish to extend this impatience across generations.[25] To some extent, as noted in Section 2.1, the average individual's investment behavior appears to reflect his own mortality risk, in combination with imperfect "intergenerational altruism". That is, people on average spend as if they do not care about their heirs' welfare (at least through consumption) as much as about their own.[26] The *total* time preference rate reflected in the interest rate is thus still higher, for such a philanthropist, than the time preference rate she would endorse. And as we look further into the future, comparing benefits from one generation to the next, this philanthropist's time preference rate falls to zero, while governments, firms, and typical households organize their activity around a principle of discounting future wellbeing at a positive rate indefinitely.

2.5 *What an efficient but impatient world underfunds*

The allocation of capital can be both efficient and impatient. That is to say, even if capital were allocated such that there are no "free lunches"—no ways to make everyone better off that aren't implemented because of a lack of perfect information or perfect contract enforceability—capital might still be improperly allocated from a patient perspective. To distinguish the philanthropic implications of impatience from those of inefficiency, and to explore how these implications interact, let us begin by considering what is most underfunded (from a patient philanthropist's perspective) in an efficient but impatient world.

In myriad ways, capital will be allocated differently in an efficient but impatient world to how it would be allocated in an efficient world in which individuals typically exhibited no time preference. Some things will be overfunded, from a patient perspective, and others underfunded. A patient philanthropist's project is then, roughly, to identify what areas will be

[24] See Freitas-Groff (2020) for a recent experimental illustration of this phenomenon and review of the literature.

[25] This pair of positions was endorsed by the Council of Economic Advisers to President Obama (2017), for instance.

[26] See Lockwood (2012) for a review of relevant evidence.

most underfunded from her perspective and then to act strategically to increase funding in these areas as much as possible.

The most straightforward example of a good that will be underfunded from a patient perspective is future consumption. Out of impatience and a lack of intergenerational altruism, people consume too much and save too little for themselves or their descendants to consume in the future. As we have already seen, this implies that a patient philanthropist would generally do better to invest for the sake of funding the consumption of his beneficiaries or their heirs in the future than to fund their consumption today (at least if he could simply "balance the portfolio", and did not have to tackle the game-theoretic considerations discussed briefly in Section 1 and at more length in Sections 3 and 4).

Many projects which require funding today are likewise underfunded from a patient perspective, because most of their benefits are expected to accrue in the relatively distant future. Projects to reduce the risk of catastrophes that would impose long-lasting damages, as we saw earlier with regard to asteroid monitoring, constitute one particularly important category of these. Another category includes projects with a chance of producing long-lasting benefits, such as many kinds of scientific research.

However underfunded these future-focused projects may be, however, *investment for the sake of analogous projects in the future* is likely to be more underfunded still.

Consider the asteroid monitoring example again. As a society, we have explicitly decided to dedicate some of our wealth to monitoring efforts today. We have also implicitly decided to leave some wealth untaxed and unspent, to be reinvested across the generations, so that some of it can be taxed and spent on the catastrophic risk mitigation efforts available—including continued (and presumably improved) monitoring, and perhaps deflection readiness—in the next century. Given the current allocation of spending across time, we should expect that the benefits of these two uses of wealth are currently about equal *from an impatient perspective*; if people preferred or dispreferred spending more today to investing to spend more on the projects of the future, it would be straightforward for us and our policymakers to spend more or less quickly. Asteroid monitoring is underfunded because its expected benefits are backloaded: today's asteroid monitoring accrues expected benefits from now to the day that life on earth would have perished by some other means. But investment for the sake of improved future asteroid protection in, say, 100 years produces benefits even more backloaded. It does *nothing* to protect us; *all* its benefits accrue to future generations.

Scientific research projects, and efforts to improve scientific institutions, have the same feature. Our impatient society likely underfunds them, because their benefits are backloaded, in that they accrue both to ourselves in the relatively near future and to future generations. But investment for the sake of future spurs to scientific progress are likely even more underfunded—not in absolute terms, but relative to what the interest rate allows for—because the benefits of such investment are even more backloaded.

Investment for the sake of future spending, in some area, does not necessarily produce more backloaded benefits than immediate spending in that area. Suppose for instance that there are only two opportunities to mitigate the damages of climate change: (a) funding a green technology research project today that will yield welfare benefits in 100 years (but no other benefits), and (b) investing to fund a flood prevention scheme in 20 years that will yield welfare benefits contemporaneously, i.e. in 20 years (but no other benefits). An impatient society will

relatively overfund (b). A patient philanthropist will thus prefer to allocate additional resources to (a), at least relative to (b).

The point I mean to emphasize is simply that having long-run benefits—even benefits that *grow over time*, as deflecting an asteroid or accelerating scientific progress may have—is not what an intervention needs for a patient philanthropist to prefer funding it immediately over investing to fund it in the future. What the present-day intervention needs is benefits *more backloaded* than those that will be offered by any analogous future intervention. Given sufficiently strong "path dependence", as in the example of the previous paragraph, such circumstances are conceivable. They will be discussed briefly in Sections 3.5–3.6. I will center most of the discussion of Sections 3 and 4 on what I believe is the more relevant case, however, in which further-future funding produces, "on average", further-future benefits.

2.6 *The limited implications of inefficiency*

The previous subsection presents an argument for skepticism about claims that it is best, from a patient perspective, to fund projects today. As noted, this argument is predicated on the assumption of a sort of "efficient market" in potential philanthropic projects. In brief, the argument goes, even if a project today really does offer substantial long-term benefits, it probably also offers enough near-term benefits that it will not be the *most* underfunded potential project in an efficient, impatient world. The most underfunded projects are those that will themselves arise in the future.

It is sometimes argued that present projects worth funding from a patient perspective do arise when such efficiency is absent. In particular, one often hears that a given intervention "offers returns" far higher than those offered by conventional financial investments. Maintaining a school in Uganda, for instance, might be found to offer an *internal rate of return* of 20%. That is, it might increase its students' lifetime incomes by so much that, if you allow the school to close, you would have to earn 20% per year "in the bank" with the school maintenance funds in order to offer the children the same stream of increases to their future income. (For simplicity, let us assume that educating someone only affects her income and that of her family and line of heirs. Let us ignore, for instance, the possibility that more educated citizens affect others' incomes by participating more effectively in the democratic process.)

Because global interest rates on investments with comparable risk profiles are far below 20%, this school would testify to an inefficient market. If students could borrow against their future income and were aware of the income gains to education, or if the region's governments were funding education according to the cost-benefit framework of Section 2.3, education in the region would be funded until its internal rate of return had fallen to the relevant interest rate. Under such circumstances, a party with typical time preferences would be indifferent between funding schooling immediately and investing to fund future schooling, and as in Section 2.5, we can infer that a patient philanthropist would probably prefer the latter. In the presence of an inefficiency like the one described above, however, a philanthropist with typical time preferences strictly prefers the immediate funding. For a patient philanthropist, therefore, the decision may seem ambiguous.

Even under inefficiency, however, investing *far enough* into the future generally remains preferable for a patient philanthropist to spending on an inefficiently underfunded intervention today, even if he is only slightly more patient than the norm. This is because the outsize benefits

to funding today stop growing as soon as these benefits have been acquired by impatient beneficiaries, whereas the benefits to waiting before spending compound as long as interest rates continue to incorporate a rate of time preference higher than the philanthropist's.

Consider the Ugandan school again. It will raise its students' incomes impressively. For a school's "above-market rate of return" to persist, however, its students would have to maintain consumption levels close to what they would have faced without the education, and use the bulk of their extra incomes to charitably fund schools with similar outsize returns. The students whose educations they fund would then have to do the same, and so on, indefinitely. This is entirely unrealistic.

Instead, in the most optimistic plausible case, the alumni collectively split their newfound resources between consumption and investment such that the income gains compound, through their lives and across generations, at a rate that is, in the long run, strictly below the prevailing interest rate. Even if an alumna's income gains are *all* invested—not given to charity schools, but conventionally invested, in ways that return the interest to herself—they will grow more quickly than the interest rate only as long as she continues to face private investment opportunities which (a) enjoy above-market returns and (b) cannot be funded by borrowing, due to some sort of market inefficiency. After growing her income to the point of satiating these opportunities, further gains will grow at the market interest rate if they are exclusively invested, and will grow more slowly than the market interest rate given that they are split between consumption and investment. Finally, observe that consumption gains cannot grow more quickly, in the long run, than income gains. In present value terms, therefore—i.e. after discounting future consumption benefits at the interest rate—the total future consumption gains accruing from the educational intervention will be finite.

By the same reasoning, cash transfers in the most optimistic plausible case produce a stream of consumption benefits that grows more slowly than the interest rate (and whose present value equals the size of the cash transfer). For simplicity, therefore, and to make the case for giving sooner as strong as possible, suppose we knew of an intervention that *immediately* left its beneficiary richer by some multiple—10x, say—of what it cost to fund. This would be like a school that increased the lifetime incomes of students and their heirs by $10 (in present value terms) per $1 spent on it, and immediately put these gains in the students' pockets.[27] Such an intervention would do 10x more good than an immediate cash transfer, regardless of what time preference rate is used to make the comparison.

In contrast, by the reasoning of the previous two subsections, investing for a future cash transfer does, say, 3% more good each year, as long as the beneficiaries persist in employing a 3% time preference rate. (As before, we are here putting the most interesting game-theoretic considerations aside by assuming that beneficiaries do not adjust their spending behavior in light of the coming transfer.) We might believe, "pessimistically", that the future contains no inefficiencies like today's potential school, so that the best spending opportunity in the future is a cash transfer. Even so, a patient philanthropist just has to wait 77 years before the future cash transfer does more good than the present classroom. This is thus a case, as promised in Section

[27] This is essentially the calculation the Copenhagen Consensus Center uses to quantify the benefits of various interventions, including educational interventions in the developing world (see e.g. Copenhagen Consensus Center, 2015).

1.2, in which giving slightly later helps less than giving now, but giving much later helps more. The tortoise beats the hare.

3. Patient giving

As noted at the beginning of Section 2, a philanthropist's preferences may differ from those of other actors in countless ways. Compared with the US Congress, for example, an impartial philanthropist is likely to be not only more patient, but also more concerned with the welfare of disenfranchised Americans, people in other countries, or animals. All these preference differences, with all other actors, pose game-theoretic complications. This document is not intended to offer an all-things-considered guide to spending strategy, but to shed light on one of the most important and under-considered dynamics. In particular, it is intended to argue for the importance of differences in patience and then to explore the implications of these differences. The last section did the former; this section and the next will do the latter.

By the discussion of Section 2 alone, one might be misled into believing that the logic of comparative patience produces some sort of Ponzi scheme. If giving next year always does 3% more good than giving this year, then it may seem as though it is never time to give.

This misunderstanding might arise because the argument of Section 2.5 is, in effect, a dramatic simplification even of AR. That is, though we might have thought that introducing impatient co-funders would complicate our disbursement scheduling problem, there is a sense in which it makes it easier. Instead of trying to estimate all the relevant variables ourselves—the interest rate, the rate at which it is getting more expensive to do good, and so on—we only have to consider one: the gap between our time preference rate and that of the other people funding projects like those we wish to fund. If we find that the others are less patient than we are, if we trust that they are allocating their resources efficiently from an impatient perspective in light of all the considerations at play, and if we observe that they are splitting their wealth between funding today and investing to fund in the future, then we can infer that we must prefer investing. There is only a single rate to add, and it is always positive.

By removing the assumption that the impatient world allocates its resources efficiently, Section 2.6 amounts to the move to TM. This is because, in Section 2.6, we allow for the possibility of inefficiencies that temporarily allow us to generate abnormal amounts of value. (We could call an opportunity to generate abnormal value, in this way, a "giving multiplier"). This move leaves us a second rate to consider: the rate at which the "giving multiplier" produced by opportunities to solve remaining inefficiencies is falling. If this rate is larger than the difference in time preference rates, it is better to give today than tomorrow. But the giving multiplier cannot fall at a constant rate forever, because it cannot fall below 1; a philanthropist can always simply transfer his wealth.[28] In the long run, in what is (from the perspective of opportunities for philanthropic impact) the worst-case scenario, there is again only one rate to consider, and it always directs us to shift our spending ever further into the future.

[28] *Ex post*, of course, much philanthropy may be counterproductive. From the philanthropist's own perspective, however, the worst available giving opportunity will never do less good, in expected value, than can be achieved by transferring his wealth to a beneficiary.

Both sections ask how a patient philanthropist can best allocate a *small increment* of capital. As noted in Section 1.4, however, when a non-negligible quantity of capital is to be allocated, we must account for the diminishing returns to spending in each period. When we do, we will generally conclude that it is best to spread spending across multiple periods. In particular, we will find that it is indeed best to begin spending "before the end of time".

The material for Sections 3.1–3.3 is taken from Trammell (2021). As noted in Section 2.5, I will not assume that the benefits to spending at a time all accrue at that time, but I will typically assume a monotonicity whereby further-future spending produces further-future benefits. I will briefly consider the implications of relaxing this assumption in Section 3.6.

3.1 *Patient giving with non-strategic impatient co-funders*

Let us begin by considering the case in which other parties funding the areas we also wish to support are non-strategic, meaning that their spending schedules will not respond to our behavior.

For all actors, a positive interest rate offers a reason to spend more in future periods, all else equal, and the fact of diminishing returns offers a reason to spread spending across time instead of investing everything for the indefinite future. The less patient an actor is, however, the less the impact to spending in a future period has to diminish before she is indifferent about reallocating capital across periods. The impatient-optimal spending schedule for a given budget will therefore involve more immediate spending, but slower spending growth, than the patient-optimal spending schedule.

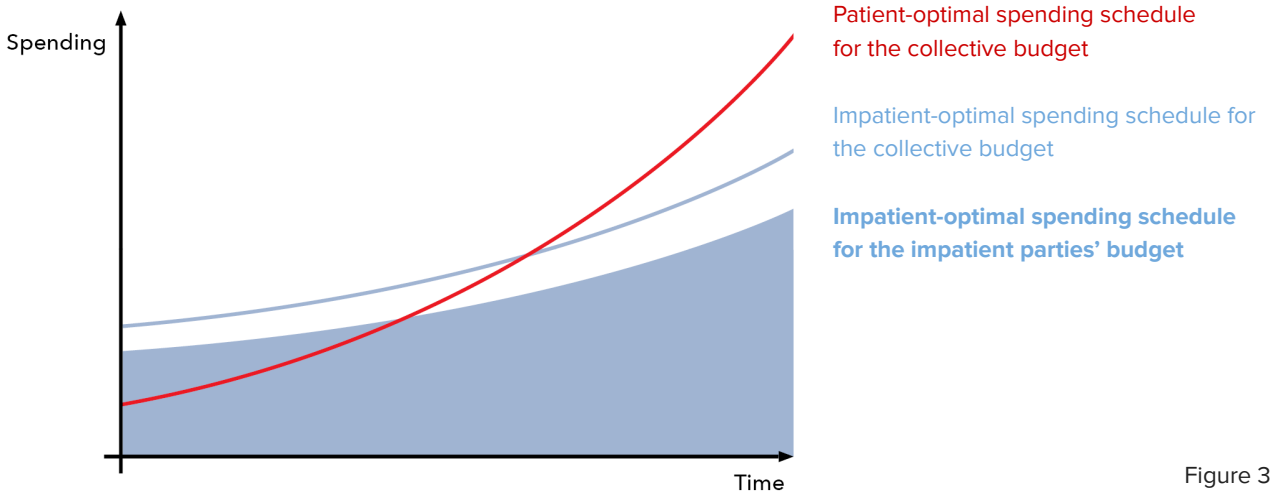

Figure 3

Impatient parties will allocate their budgets so that spending follows the path given by the pale blue region in the figure above. ("On average", across all potential funding areas, the impatient-optimal spending growth rate should roughly equal the economic growth rate in a predominantly impatient world. This will also be the spending growth rate observed if each generation spends a constant share of its income contemporaneously funding the area in question, even if no resources are explicitly invested for the sake of future funding.) If impatient parties could allocate the budgets of the patient as well, they would spend more at each period, so that spending followed the path given by the pale blue line.

From a patient perspective, however, it would be optimal to allocate the collective budget so that spending followed the path given by the red line. The red line grows more quickly, at something closer to the interest rate. The gap between the red line and the pale blue region thus grows over time, even in proportional terms, and the "underfundedness" of a given time from a patient perspective grows accordingly.

This is why an infinitesimal increment of capital does ever more good, from a patient perspective, invested ever further into the future. A patient philanthropist with a budget of any positive size, however, will pull the collective spending schedule as far as possible toward the patient-optimal spending schedule—i.e. will best "balance the portfolio"—by starting to spend at the time, marked "*t**" below, when he can *subsequently* render the collective spending schedule patient-optimal.

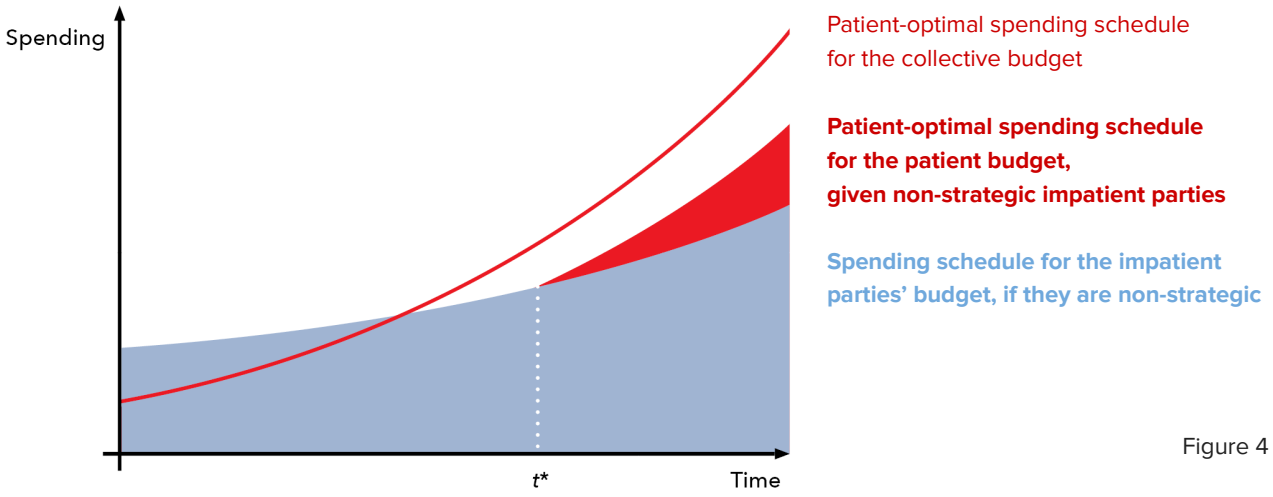

Figure 4

If the patient funder(s) in this area are sufficiently wealthy in comparison to the (non-strategic) impatient funder(s), it may well be possible to implement the patient-optimal spending schedule for the collective budget immediately. That is, a patient funder may find that the impatient-optimal spending rate of the *impatient* budget is less than the patient-optimal spending rate of the *collective* budget. This is effectively the condition the patient philanthropist finds himself in from *t** onward in the example above. Patient philanthropists will then start spending immediately, despite the presence of less patient co-funders. Their own spending, however, will still start smaller and grow more quickly than it would if the impatient parties were not present.

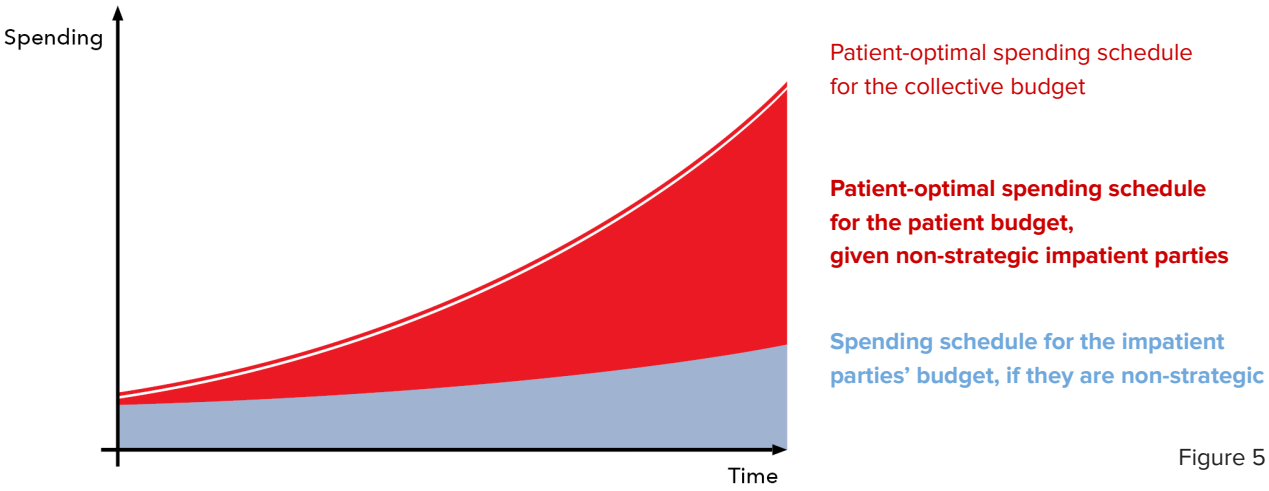

Figure 5

Finally, as illustrated, a philanthropist's optimal spending rate in a given area at a given time depends greatly on the total budgets allocated to the given area by patient and impatient parties. Note however that these total budgets include not only the budgets observed today, but also (the present values of) the budgets that will be allocated to the area in the future. For instance, even if today's patient funders are poor relative to today's impatient funders, they may find it worthwhile to spend immediately if they believe that there will be a large increase in patient philanthropic funding in the future.

3.2 *Patient giving with strategic impatient co-funders*

Now let us suppose the impatient funders are strategic. That is, just as the patient try to set a spending policy that responds optimally to that of the impatient, let us suppose that the impatient try to set a spending policy that responds optimally to that of the patient. The parties are then playing a kind of game.

If all parties could *commit* to their disbursement schedules at once, the schedules would be best responses to each other if and only if they took the following form.

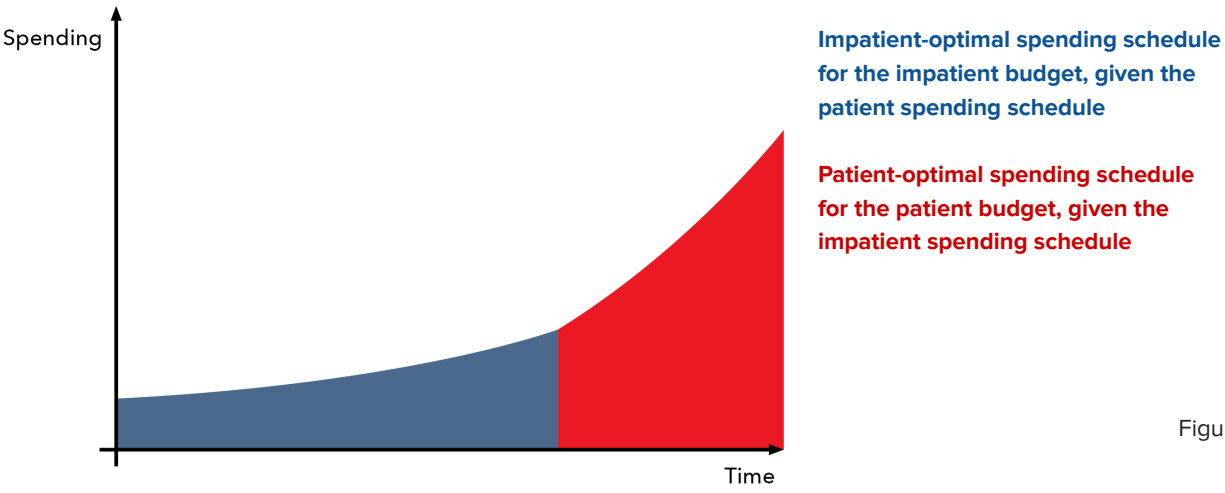

Figure 6

Note the kink. Given the overall spending growth rate before that point, impatient parties would be indifferent to reallocating capital across periods at which they do spend. They would strictly prefer not to invest additional capital for spending after the kink, considering subsequent future "overfunded" by the patient. Likewise, patient parties would be indifferent to reallocating capital across periods at which they spend, and they would strictly prefer not to spend anything before the kink, considering the short term overfunded by the impatient.

In practice, however, funders generally cannot commit to a given spending schedule. A funder's spending rate at a time, once it arrives, will depend on how much each party has already spent and on how he believes others will adjust their future spending in light of his own immediate spending.

This fact gives rise to multiple "equilibria". That is, there are multiple ways that the impatient parties could set their spending policy, and the patient parties could set theirs, such that what everyone is doing is a best response, by her lights, to what everyone else is doing. Without knowing which equilibrium we find ourselves in, therefore, there is simply no way to define the patient-optimal disbursement schedule.

One equilibrium, however, is much simpler than the rest. It is therefore what we might expect to observe among strategic funders if they do not agree to coordinate in some other way. This is the equilibrium in which the patient parties withhold all spending as long as they observe the impatient parties "overspending", and in which impatient parties know that patient parties will do this. In this equilibrium, the spending schedules take the following form.

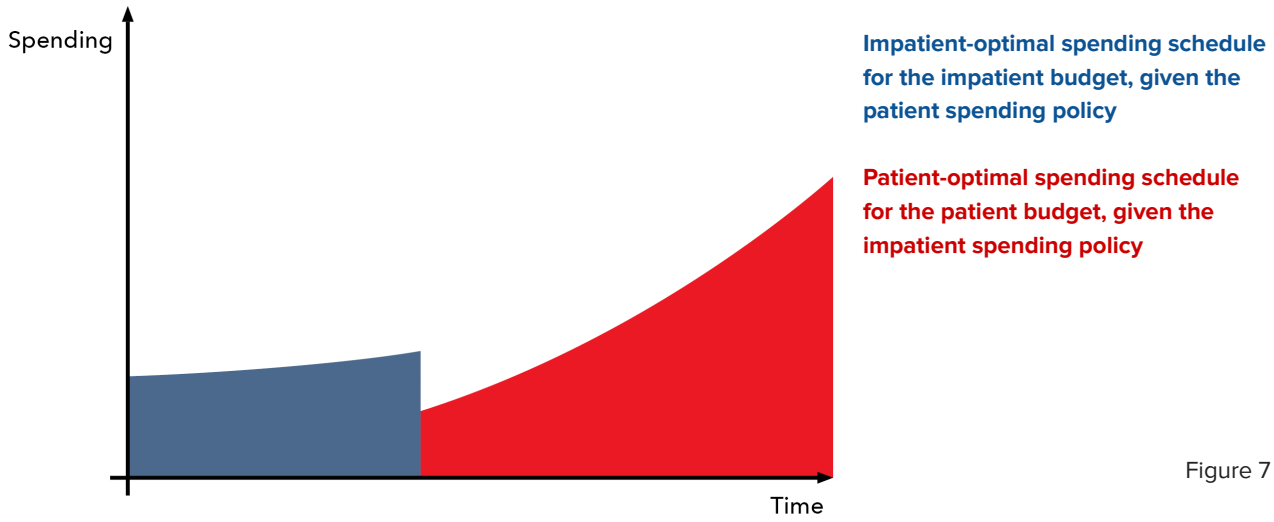

Figure 7

As in the case in which both sets of parties commit to their schedules in advance, each group here allocates its budget so that it is indifferent to reallocating capital across the periods at which it spends. This implies slow growth during the "impatient regime" and fast growth during the "patient regime". Here, however, the impatient can take advantage of the patient parties' lack of commitment power by spending down their resources more quickly. Once they have no more to spend, the patient have no rational option but to begin disbursing their own budgets patient-optimally. The impatient thus effectively shift resources to the near term—partly at the

expense of the medium term, to their moderate disappointment, but also at the expense of the long term, which they do not mind.

The resulting spending schedule is the same as would obtain (in a "simple equilibrium") if the impatient had the power to commit to a disbursement schedule but the patient did not. There is thus a sense in which the impatient parties in this game have a "first-mover advantage".

The collective spending schedule of Figure 7 is *inefficient*. The spending rate falls sharply at the regime-switching time, so, because of the diminishing returns to spending at each time, the value of spending just after the regime-switch rises sharply. The impatient funders would therefore prefer to shift resources from just before the regime-switching time to just after. If the patient could agree not to respond to this spending delay by withholding their own spending further, the resulting smoother spending schedule would be preferable to all parties.[29]

An efficient collective spending schedule would take a shape like the below: it would exhibit a growth rate that begins somewhere between the impatient- and patient-optimal growth rates and rises smoothly over time to the patient-optimal growth rate. If the starting growth rate is not too low, the patient will prefer it to the Figure 7 collective schedule, and if it is not too high, the impatient will as well. Moreover, if the parties can observe each other's behavior, they do not need to sign an enforceable contract to implement this collective schedule. As long as they all follow the spending schedules below, and adopt the policy that if anyone deviates (i.e. if the patient try spending too little or the impatient too much) then they will subsequently follow the Figure 7 schedules, no one will at any point want to deviate. This is thus one of the "less simple" equilibria referred to above.

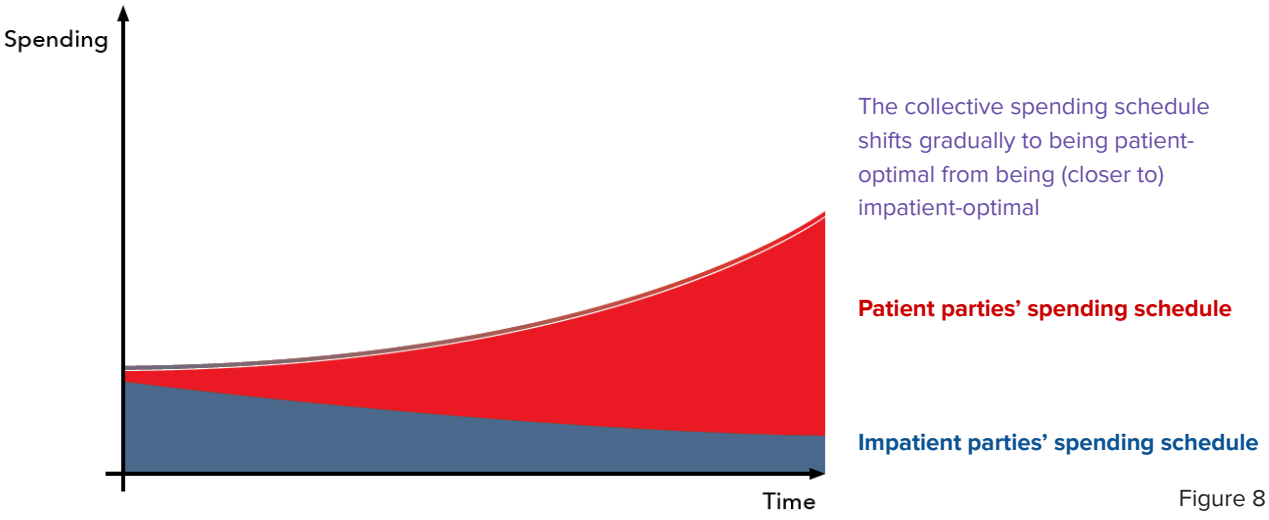

Figure 8

As explained at the beginning of the previous subsection, the more patient a funder is, the slower he would prefer the collective spending schedule to begin and the more quickly he would

[29] More subtly (and less relevantly for our purposes), the spending schedule of Figure 6 is inefficient as well. The impatient are indifferent about small reallocations of their own spending forward, toward the kink, but the patient strictly prefer them. Likewise the patient are indifferent about small reallocations of their resources backward, but the impatient strictly prefer them. All parties would therefore prefer an agreement whereby some resources were shifted toward the kink from both directions.

prefer it to grow. When funding an area for which he is the *only* funder, therefore, he will tend to disburse a smaller fraction of his budget each period than the budget-fractions impatient parties tend to disburse each period in the domains for which *they* are the only funders. (This is only a rule of thumb, since atypically fast spending may be warranted in certain areas, including areas funded only by the patient.)

It may be worth noting a pattern that appears throughout these past two subsections: that the presence of impatient parties funding a given area induces (strategic) patient parties to spend even more slowly than they would otherwise.[30] When facing comparatively well-endowed but non-strategic impatient parties, or when engaging with strategic impatient parties of any size in a "simple" equilibrium, the patient do best to begin by spending nothing. When facing comparatively poorly-endowed and non-strategic impatient parties, or strategic impatient parties of any size in a "complex" but efficient equilibrium, the patient do best to begin by spending at a positive rate, but still more slowly than the (usually) slow rate that would be optimal if they were alone.

3.3 *The gains to giving patiently*

Throughout this section, we have discussed the shape of the *optimal* spending schedule for patient givers in various circumstances. The magnitude of the benefit to employing this schedule, instead of "following the crowd" and disbursing on an impatient-optimal schedule, is a further question.

If a patient actor intends to spend all his wealth in an area for which he is (and will always be) the only funder, the gains to spending patient- rather than impatient-optimally depend entirely on the details. In principle, these gains can be arbitrarily large. In practice, if the relevant variables—such as the extent to which spending at each time faces diminishing returns—are similar for the funding area in question as for spending on individual consumption, then the gains from moving to the spending schedule that is optimal given perfect patience, from the schedule that is optimal given a typical rate of time preference, appear to be relatively modest. I estimate that the gains from this switch are roughly equal to the gains from staying on the impatient schedule but having a budget that is 30% larger.

Suppose that the area in question is also funded by impatient parties, however, and suppose that the impatient parties are non-strategic. In this case, for the reasons given in Section 2.5, each year's wait allows a comparatively small patient party to do a bit more good, from his perspective—roughly 3% more, typically, if the gap between the parties' time preference rates is 3%. As noted in Section 3.1, this relationship persists until the patient party's budget has grown to a large enough share of the collective budget that, for reasons of diminishing returns, his further investment would substantially decrease the impact of further future spending.

The smaller the budget of the patient is relative to that of the impatient, the longer the patient will be able to accrue these time-preference-based benefits by waiting to spend, before the diminishing returns effects grow substantial. This allows us to draw the following

[30] The presence of patient parties likewise induces strategic impatient parties to spend even more quickly. The more general phenomenon exhibited here is sometimes called "strategic polarization", and is explored by Kalai and Kalai (2001) and many others.

conclusion: that as the share of the initial collective budget in patient hands goes to zero, the extent to which he "does more good" by following the patient-optimal spending strategy than by spending impatient-optimally giving grows infinitely large. A small patient funder in a large impatient world should be willing to spend approximately his *entire* budget, if this were necessary to secure him the right to spend the remaining pittance patiently.

The same result can be shown to hold for small patient parties facing strategic impatient parties, and engaging them in either the "simple equilibrium" or a "complex but efficient equilibrium".

Let us illustrate these results with two very rough attempts to quantify the gains from patience available in practice.

First, suppose a total of $20b were allocated today to the cause of direct spending, in the patient-optimal way, on consumption goods worldwide. This is of course the broad area which almost all the world's wealth is already directed to funding, almost all of it impatiently.

If impatient parties are fully non-strategic, in that they do not consume their own wealth more quickly in light of the coming flood of patient philanthropy, then, at least given what I believe are relatively conservative estimates of the values of the relevant parameters,[31] the patient will do best to invest all their wealth for 424 years before subsequently implementing the patient-optimal collective spending schedule. Doing this will, on these estimates, do 580x as much good from a patient perspective as would be done by spending the $20b on typical consumption goods impatiently.

If impatient parties are fully strategic, optimal patient behavior in the "simple equilibrium" entails investing for 437 years before spending, and does 291x as much good from a patient perspective as done by acting impatiently. Optimal patient behavior in a "complex but efficient equilibrium" could entail beginning to spend a small amount immediately, but would do *at least* 291x more good from a patient perspective than acting impatiently.

There are many obstacles to taking disbursement plans on this scale literally. If patient philanthropists tried to invest until their assets constituted a large share of all global wealth, this project would likely have a host of economic and political ramifications we have so far managed to ignore. Some of these complications are discussed in Section 5. Nevertheless, the exercise above demonstrates that patient philanthropic actors in an impatient world may do well to invest most or all their funds for the relatively long run, and that spending impatiently may constitute a highly significant error. If the details of the calibration as presented here fail, that is because the qualitative argument is a victim of its own success, recommending investment beyond the scale that simpler models of disbursement scheduling are even equipped to consider.

Second, by contrast, let us briefly consider a hypothetical funding area for which the patient and impatient budgets are initially equal. If the relevant variables (such as the rate of diminishing returns) are the same in this area as in the area of direct spending on individual consumption, the patient should begin spending immediately, given non-strategic impatient funders. Given strategic impatient funders, the patient should invest all their wealth for no more than 30 years. In either case they will do only about 1.4x as much good, from a patient perspective, by following the patient-optimal policy as they would by acting impatiently.

[31] I assume a time preference gap of only 1.5% per year, for instance, rather than the 3% used in Section 2 for illustration. See Section 5.1 for more information about the choice of time preference gap.

3.4 *Time and space*

Sections 3.1–3.3 focus on what might be called *crowding out across time*. The presence of impatient parties of one kind or another should induce patient parties to cut back on near-term spending, and the presence of patient parties will likewise induce strategic impatient parties to cut back on investment for long-term spending.

Another important dynamic, as perhaps illustrated in the Gates Foundation example of Section 1.5, is *crowding out across funding areas*. Because some domains of potential funding systematically offer more backloaded benefits than others, they are more underfunded from a patient perspective in an impatient world. Patient philanthropists will want to devote larger shares of their budgets to these domains. Put another way, the relative overfunding of projects with short-term benefits should (and clearly does) crowd patient philanthropists out of those and into other projects. Likewise, the concentration of patient philanthropic funding on projects with backloaded benefits will crowd strategic impatient funders in the other direction.

A particularly unambiguous (if self-serving) example of a project more valuable to patient than impatient actors is research on patient funding strategy. Indeed, because the aim of this research is to shift the global distribution of spending toward projects with long-term benefits and away from those with short-term benefits, much of it is *dis*valuable from an impatient perspective. Unsurprisingly, all such work, including this specimen, appears currently to be funded by patient philanthropists.

Beyond investigation into niche strategic questions for patient actors, perhaps the quintessential example of an area promising primarily long-term benefits is existential risk reduction. Assuming that the wellbeing produced by human civilization is positive and growing, and that it will continue to grow for the foreseeable future, averting an existential catastrophe produces a stream of benefits that lasts as long as civilization itself and grows over time as civilization progresses. Furthermore, even if spending on risk reduction infrastructure does at some point avert an existential catastrophe, it is likely to do so long after the infrastructure is first established. As noted in Section 2.3, the result of all this backloading is that large, impatient governments do fund existential risk reduction efforts to some extent, but not nearly as generously as would be justified from a more patient perspective. A substantial fraction of all global funding on explicit efforts to reduce existential risk appears to come from philanthropic actors with particular concern for the long-term future.[32]

Another broad field in which the benefits to spending often accrue atypically long after the funds are spent is basic scientific research. This is because basic research often takes many years to yield more efficient production techniques or new varieties of final goods, and more

[32] The fraction is of course highly sensitive to what is counted as an effort to reduce existential risk. According to at least some researchers of existential risk, such as Ord (2020), the two most likely sources of existential risk are extreme pandemics and misaligned AI, which together account for more than half the total risk this century. Todd (2017) estimates that global annual expenditures on extreme pandemic prevention and AI safety research respectively totaled $1b and $10m at the time, and Todd (2021) estimates that annual expenditures on biosecurity and AI risk mitigation from within the Effective Altruism community alone totaled $41m and $40m in 2019. These categories may not perfectly align across the two periods, but they do testify to the recent growth of the field of AI safety and the extent to which its funding comes overwhelmingly from a small number of philanthropists.

importantly because, at least under some circumstances, the economic growth eventually produced by accelerating basic research can compound into the subsequent future. Because of the extent to which further consumption increases contribute less to people's wellbeing as they get richer, however, even eternally compounding boosts to the growth path of consumption per capita probably produce welfare benefits less backloaded than the welfare benefits of existential risk reduction. Suppose, for illustration, that the relationship between consumption and individual welfare is *logarithmic*.[33] A growth boost that compounds forever at a constant rate then produces a *constant*, rather than growing, stream of welfare benefits. At least in theory, therefore, we should expect basic research to be less underfunded from a patient perspective.[34] In practice, while many philanthropists do fund basic research out of concern for the future, governments remain the area's largest funders.[35]

By contrast, the benefits to spending on immediate consumption are entirely frontloaded. Patient philanthropists should thus be relatively uninterested in the area. Still, though, impatient parties may sufficiently underinvest for the sake of consumption far in the future that patient philanthropists would do best to fund far-future consumption to some extent, rather than simply spending as much as possible in the domains above.

Finally, of course, some classes of expenditures, such as recreational drug consumption, produce short-term benefits but greater longer-term costs. These will presumably be funded exclusively by (especially) impatient consumers, and patient philanthropists will try to avoid funding them on any timescale.

Let us conclude by considering crowd-out across both time and funding area, at least for the areas listed above, given the existence of both strategic and non-strategic impatient funders. Modeling this problem realistically and quantitatively would be difficult, especially in light of interactions across areas, such as the many possible relationships between scientific research and existential risk. But if patient parties act strategically, and if they cannot coordinate (or

---

[33] Even this would generally be considered optimistic: while some ways of estimating the relationship between consumption and welfare conclude that the relationship is roughly logarithmic, others conclude that the benefits of consumption increases diminish more sharply. In support of the logarithmic view, see Chetty (2006) on the risks people are willing to take in the context of the labor market, Gandelman and Hernández-Murillo (2015) for global survey data on the relationship between consumption and wellbeing, or Drupp et al. (2018) for a survey of 173 domain-expert economists (whose median response implies a logarithmic relationship, but with a mean response implying a more sharply diminishing relationship).

[34] This reasoning assumes that the benefits of economic growth consist primarily of increases to future consumption per capita. To the extent that the benefits consist of increases to future population size, we reach a similar conclusion: even if accelerating growth hastens the day when human (or posthuman) civilization begins expanding seriously across space, so that long-run population growth—and the population benefits of averting existential catastrophe—is cubic, the long-run population benefits of accelerating growth are only quadratic. (A more comprehensive suite of analyses comparing the benefit-streams produced by growth acceleration and existential risk reduction is given by Ord (in progress).) If economic growth also substantially lowers existential risk, however (as proposed e.g. by Aschenbrenner (2020)), the relative backloadedness of the benefits produced by the two areas is ambiguous.

[35] Most OECD countries spend well over 0.25% of GDP on basic research (OECD, 2020), implying annual expenditures of at least $150b from the OECD alone. Funding for R&D more broadly is over an order of magnitude higher (UNESCO, 2020). This is much greater than annual philanthropic research spending, a point emphasized for instance by Callahan (2018) and the research-funding philanthropists he interviews.

contract) with the world's diffuse array of governments, firms, and other impatient funders, the respective disbursement schedules will, I would guess, look qualitatively like the following.

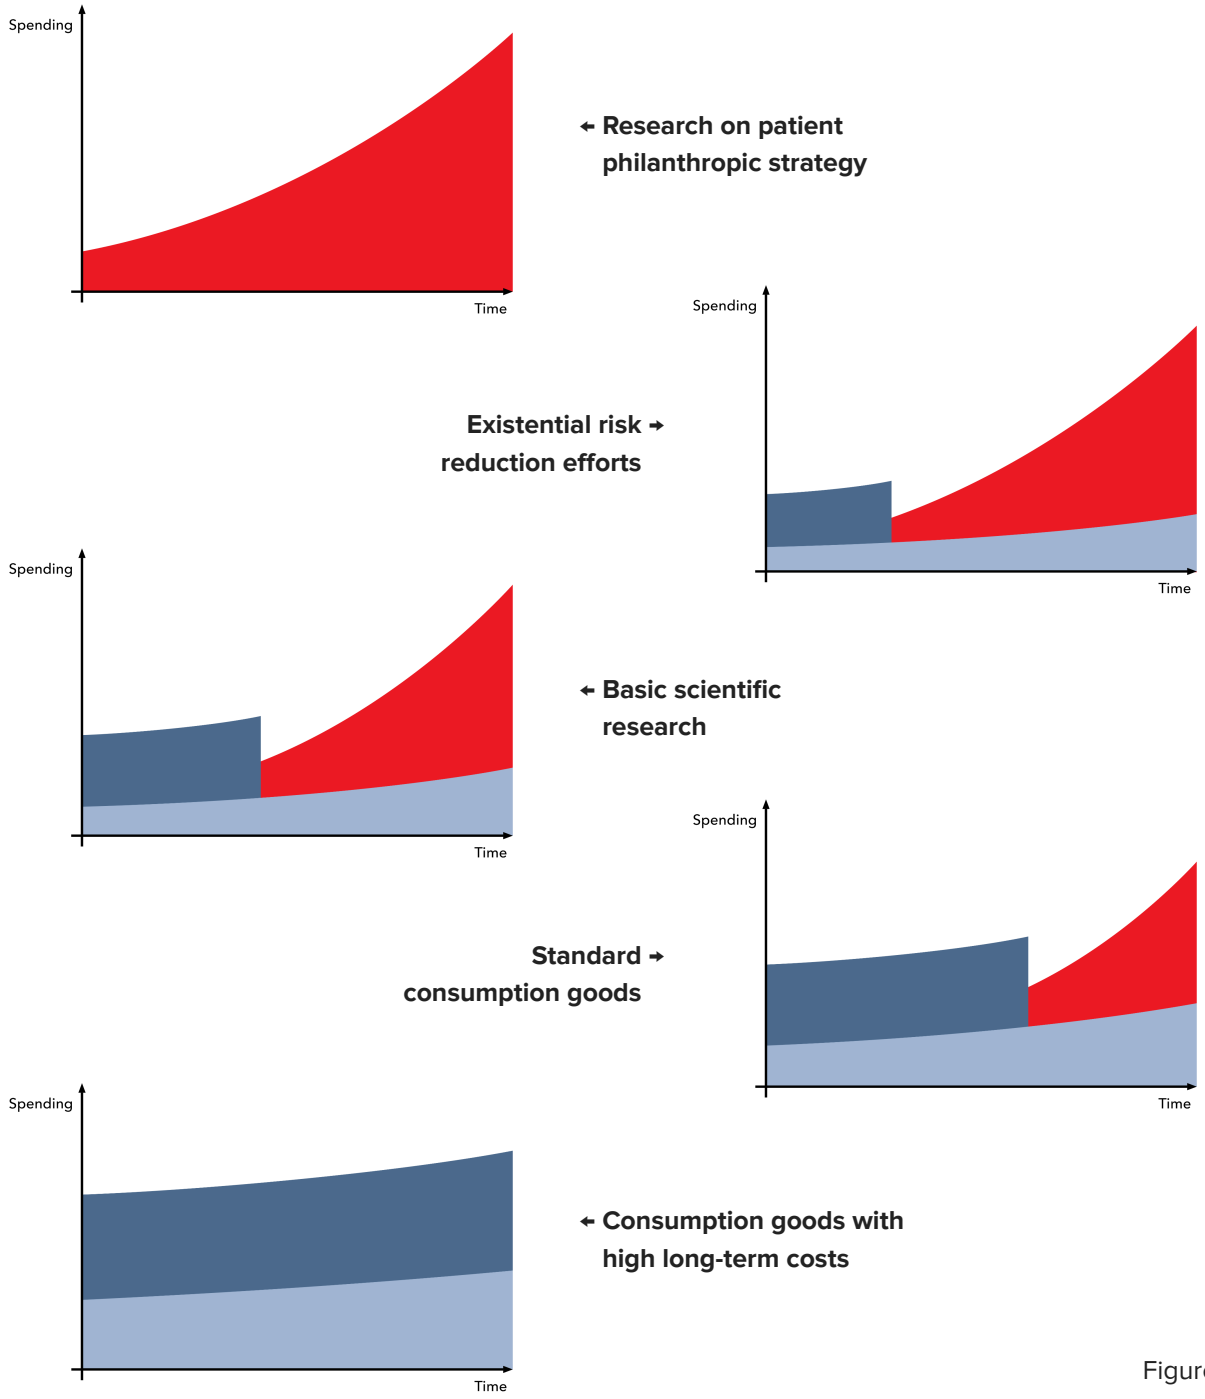

Figure 9

Once again, we see a channel through which patient parties' disbursement should be expected to begin more slowly, and grow more quickly, than it would in a patient world. As time

proceeds, it is optimal not only to increase spending in a given area but to expand the scope of patient philanthropy.

3.5 *Uncertainty and change*

Many of the variables relevant to disbursement scheduling are subject to substantial uncertainty, even in the present. We do not know how severely returns diminish, for example; if spending in a given area doubled immediately for one year, even if we are confident that less than twice as much good will be done in that area this year, it is rarely obvious how much less. Relevant variables may also change over time, predictably or unpredictably. We may know the current interest rate, but it is unlikely to remain at its current value forever. It, and other variables, may even change as a result of what we do, if we are sufficiently large and effective players.

It is important to account for these uncertainties and changes when setting a spending schedule in isolation, as one must do when one is the only agent funding an area. By and large, however, they should not qualitatively affect the interaction between patient and impatient funders as explored above. Whatever the future may hold for, say, interest rates, however uncertain we may be about them, and however much our actions may influence them, patient funders will still typically want to backload spending relative to impatient funders, and they will still want to skew spending toward projects with backloaded benefits.

That said, there are two classes of questions in this vein to which even the broad outlines of the patient philanthropist's spending problem, as sketched above, are likely to be sensitive.

The first concerns the arrival of future funding.

As noted at the end of Section 3.1, what we mean by the "total budget allocated to a given area" by parties of some time preference rate incorporates the present values of the budgets that will be allocated to the area by funders of the given time preference rate in the future. This point is especially important today because, as documented by Callahan (2018) and as demonstrated by the number and growth of Giving Pledge signatories, recent decades have seen a veritable explosion in American philanthropy. Many of these new philanthropists express a desire to use their wealth for the sake of future generations. If some of these funders are truly patient in the more general sense, and if growth in their ranks is expected to continue, many areas in which it would otherwise seem best for patient philanthropists currently to spend slowly, or not at all, may actually already be worthy of substantial patient philanthropic spending.

Indeed, due to borrowing constraints, it may even be optimal for today's patient philanthropists to spend more quickly than is feasible. That is, since the collective patient philanthropic budget today includes the budgets of unknown individuals who have not yet begun to engage in philanthropy—not to mention, presumably, the budgets of many more who have not been born—it may be that the optimal way of allocating these resources across time would be to borrow, spend more quickly today, and pay back the debt with the resources that will arrive in the future. Of course, this cannot be done. Under these circumstances, the next-best option for today's patient philanthropists would be to time their spending so that their endowments are entirely spent down as soon as new patient funding arrives.

As noted in Section 1.1, however, "giving while living", and spending quickly in general, are increasingly popular. The current boom in philanthropic activity, including activity on projects

with delayed benefits, may therefore be no more than a fad. In this case, the current philanthropic influx is better modeled as a temporary increase in impatient funding. This should have an effect opposite to the effect of an increase in patient funding: it should crowd patient funding *out* to the future. Also, to the extent that this is even a serious possibility, it may be particularly important for patient philanthropists to invest for the sake of "philanthropic insurance", in the event that current rates of philanthropic funding growth, especially in the most important areas, do not last.

This observation suggests a potentially neglected area of philanthropic strategy research: namely, research into future philanthropic preferences. Just as the terms of a 30-year loan to Coca-Cola depend in part on multi-decade forecasts of global demand for soft drinks, financial decisions in the philanthropic sphere should be responsive to long-run forecasts of other funders' spending patterns. For instance, if the founders currently passing through Y-Combinator largely reject their elders' calls for giving while living, and plan to use what wealth they earn to establish slow-disbursing, long-lasting trusts, this would be highly valuable for today's patient philanthropists to know.[36]

The second concerns the arrival of future funding *areas*.

The previous subsection considers patient funding strategy across a handful of funding areas in which we can currently spend and of which we are currently aware. But many of today's philanthropic concerns, at least if sufficiently narrowly construed, could not have been funded in the past. No medieval benefactors could have spent contemporaneously on nuclear disarmament efforts, and funding to speed scientific progress may have been inconceivable to them. Likewise, the future may uncover philanthropic funding opportunities, including some particularly valuable from a patient perspective, which are unavailable or inconceivable to us today. If plans are ever made to establish a world government, lobbying the authors of its constitution would be a new philanthropic project with potentially very long-lasting impacts.

To the extent that we believe that a large fraction (in present value terms) of all patient philanthropic spending will ultimately best be spent in these sorts of future funding areas, we should currently spend the patient philanthropic budget more slowly. We should also, perhaps, devote some effort to forecasting these future funding areas and preparing to support them effectively once they arise.

On the other hand, there may be valuable funding areas today that will eventually vanish, or shrink substantially in scope, without being replaced by similarly worthwhile funding areas in the future. An illustration of one way this may happen is suggested by the example above: it will presumably not be possible to influence the authorship of the world government's constitution for very long. Another relates to the Ugandan school discussed in Section 2.6: various investments and public goods are inefficiently undersupplied in the developing world, creating opportunities for outsize philanthropic impact, but these may vanish as the countries in question develop and their markets and governments grow more efficient.

As noted there, there is at least one funding area that cannot vanish (except under truly catastrophic circumstances), namely the direct provision of consumption goods. An increment of capital can thus likely do more good in an impatient world by serving this long-lived funding

[36] For an example of such research with respect to spending by the Effective Altruism community over the next decade, see Dillon (2021).

area in the future than by serving a seemingly more promising, but shorter-lived, funding area while it lasts. As we have seen throughout Sections 3.1–3.4, however, after recognizing that patient philanthropic resources are not mere increments but positive sums, when optimally allocated they might be able to substantially decrease the value of investing for future consumption spending. In particular, they might do so by enough that it is worth funding alternative, and more efficient, projects in the present—at least after crowding out all impatient, strategic funding in these areas. If we believe that this point has been reached, so that a given area is worth funding today, then we should spend more quickly (in all areas) after accounting for the fact that some areas will not be able to use our funding in the future than we would if we thought they would require a large funding stream forever.

3.6 *The hinge of history*

As noted in Section 3.5, funding areas of particular interest to patient philanthropists—that is, funding areas with particularly backloaded benefits—come and go. Since the boundaries between "funding areas" are vague, we might equivalently say that the extent to which funding in a given area promises backloaded benefits may rise and fall. Instead of a clean division in which impatient funders spend in the near-term and patient funders invest to spend in the future, therefore, we should expect to see a somewhat more complex spending pattern.

In a "simple" equilibrium, the result should look roughly as depicted at right. When certain funding sub-areas provide opportunities for unusually more long- term than short-term impact, patient philanthropists will wish that these were better funded.

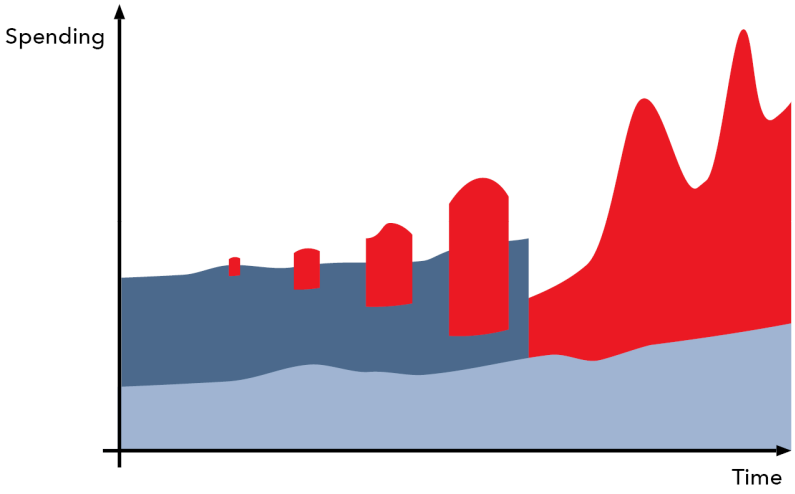

Figure 10

If this preference is extreme enough, the patient will be willing to fund these sub-areas themselves. Because their funding will crowd out impatient funding, however, they will focus their spending only on those narrow projects for which the ratio of long-term to short-term benefits is greatest, leaving the impatient to fund the rest. Early in time, few projects will meet this bar. As time goes on, however, with the impatient usually exhibiting higher spending rates than the patient, the patient grow relatively richer. Projects they are willing to fund, at the cost of crowding out impatient funding, therefore arise more frequently. In the long run, patient parties cover all funding in the area, as seen earlier in this section.

As presented in Figure 10, this dynamic is a second-order consideration. In principle, however, it may be a more important disbursement scheduling consideration than anything else we have discussed. That is, if the spikes in "potential for foreseeable, outsized long-term impact" are big

enough, the central lesson of patient philanthropic strategy may be to save for these spikes and then to disburse quickly when they arrive.

Note that this would require large spikes in *foreseeable* long-term impact. If our ability to distinguish the times with most opportunity for backloaded benefits is weak, the parties' spending schedules should look more similar to those presented in Sections 3.1–3.4. To shed light on whether this ability is weak or strong, Risi et al. (2019) analyze a corpus of almost 2 million comments by members of the US State Department on news events from 1973 to 1979. The authors find that these contemporary judgments of an event's long-term significance were only weakly correlated with judgments of the event's significance by historians with four decades of hindsight.

Nevertheless, a case can be made that, even though it is hard to judge the historical significance of a particular news event contemporaneously, it may be easier to evaluate the significance of historical eras on a longer scale.

It may be argued, for instance, that the survival (or, in some broader sense, the trajectory) of human civilization is currently at risk; that philanthropic funding can be used to significantly lower this risk; that this risk will soon pass (for better or worse); and that when it does, comparable opportunities for philanthropic spending to yield backloaded benefits are unlikely to arise within the lifetime of any philanthropic institution we establish today. Then the aggregate spending schedules of patient and impatient parties should be expected to look roughly as depicted at left. As we can see, the considerations reflected in Figure 11 are the same as those reflected in Figure 10, but their relative importance differs greatly.

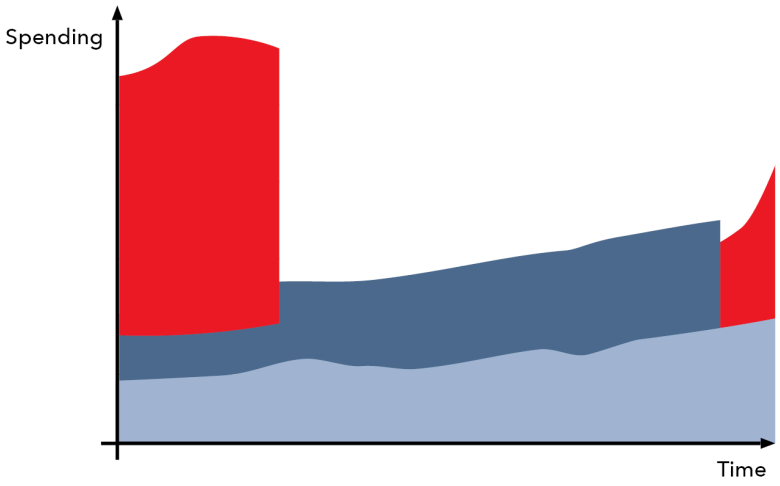

Figure 11

In fact, it is a common belief among futurists and future-oriented philanthropists that these are essentially the circumstances in which we currently find ourselves. The position was perhaps first famously articulated by Parfit (1984, p.351), who argued that as a result, our age was "the most important in human history". The claim was put more succinctly by Sagan (1994, p.306), who maintained that we were living at the "time of perils". Parfit (2011, p.616), returning to the argument, held that we were living at "the hinge of history". Ord (2020) defends the position most extensively, calling our era "the Precipice". Finally, see Karnofsky's (2021) blog series. Though the positions of these authors differ somewhat, let us collectively call their position the "hinge of history hypothesis".

If we accept this hypothesis, its implications for the timing of philanthropic spending depend greatly on how long we expect the hinge to last. If it will only last for another decade, say, then

we should quickly spend all (or almost all) the resources at our disposal, and there is little value in developing long-term investment plans.

Many who defend the hinge of history hypothesis, however, have a much longer time horizon in mind. Parfit, for instance, in both the texts cited above, characterized this important episode of history as "the next few centuries". On this view, it is again an open question whether patient philanthropists should spend anything today.

As noted in Sections 2.5 and 3.4, there are some existential risk reduction domains, such as asteroid monitoring, in which impatient parties are already spending to some extent. We should expect that in these domains, investing a small sum for future spending (while any risk remains in this domain) does more good than spending the small sum immediately.

To be sure, for some of these domains, it may be optimal for patient philanthropists to spend not a small sum, but enough to crowd out strategic impatient funding all the way back to the present. Also, there are other domains of promising existential risk reduction projects which receive no support from impatient parties at all. In these areas, as depicted in Figure 11, spending by the patient should begin at once. However, unless we expect the cost of lowering existential risk to grow very quickly over the coming centuries, the prospect of investing at interest should still lead us to skew our spending somewhat toward the end of the "hinge of history".[37] What might appear to be a countervailing reason to spend somewhat toward the beginning—namely that while we delay, we endure a higher risk rate—should itself be largely offset by the fact that higher risk rates should be expected to produce higher interest rates.[38]

Finally, note that the case for centering the patient philanthropic spending schedule around a temporary opportunity for long-term impact, as framed here, depends on an absence of similar opportunities arising later within the lifetime of a philanthropic institution existing today. As discussed at more length in Section 5.1, I believe that institutions can last for centuries, but I do not believe they can be expected to last for millennia. By contrast, Parfit, Sagan, Ord, and Karnofsky argue that we live during the most important era *of all time*.

Of course, this is a much stronger claim. MacAskill (2020) attacks it by arguing *inter alia* that, because human history has been (and, hopefully, will be) so long, our prior probability of finding ourselves in the most important era of all time should be miniscule. He proposes that, though we may have some evidence in support of the hinge of history hypothesis, the evidence is insufficient to overcome this strong prior—and that, therefore, patient philanthropists should generally invest their resources for use by future generations. This argument attracted vigorous debate on the EA Forum, including from Ord.

While MacAskill's article and its replies provide interesting food for thought on our possible place in the universe, they are largely inconsequential to the disbursement scheduling problem. Suppose we face a ten percent chance of existential catastrophe this century, which our

[37] On the other hand, uncertainties about the timelines of certain risk-mitigation opportunities can motivate allocating resources sooner, as modeled by Cotton-Barratt (2015).

[38] Note that this would not apply to all fleeting opportunities to have long-lasting impacts. For instance, suppose that settling other star systems would be an important contribution to the value of the future, and suppose that, every year, there is some risk that an irrevocable and perfectly enforceable global prohibition on leaving Earth will begin. Under these circumstances, it would be important to quickly spend patient philanthropic resources on rockets and so on. Since most investors would be unaffected by this prohibition, however, it would not affect interest rates.

philanthropy can shrink substantially, and no probability of a catastrophe for the subsequent million years. The existence of an even more cost-effective opportunity to save the world in one million years would make existential risk reduction today no less worthwhile. Likewise, if there *will* be more cost-effective existential risk mitigation opportunities over the next few centuries, spending today is probably a misuse of resources, even if we happen by a remarkable chance to be living in the most important millennium of all time.

In short, whether or not we believe we are in an era with outstanding opportunities for foreseeable long-term impact, we must consider the empirical details about how these opportunities, and others' funding on them, will unfold over the next few centuries. We cannot find the answers to this question from first principles. If we are confident enough that a period rich with future-focused funding opportunities is underway, will not last long, and will not soon return, we should quickly spend all that we have. Otherwise, we will likely still do best to invest at least some patient philanthropic resources across the generations.

4. Patient philanthropic investing

We have implicitly assumed throughout the discussion so far that there is only a single investment asset class. We have also assumed that while capital is being invested, it never does good or harm beyond what is captured in returns to the investor. When we relax these assumptions, we confront the question of how a patient philanthropist should invest. We must also determine how these investments' risks and externalities change the disbursement scheduling conclusions presented in the previous section.

In some ways, the question of how a patient philanthropist should invest reduces to the question of how anyone should invest: a complex and contentious enough question on its own, but one about which an immense amount has already been written. Here, therefore, we will focus only on the ways in which patience, philanthropic goals, or the combination of the two might motivate investment choices different from those that are optimal for investors with more typical preferences.

4.1 *Impact investment and divestment*

Whereas a self-interested investor cares only about her investment's return profile, a philanthropic investor must also account for the ways in which her investments benefit or harm others—or more precisely, achieve or frustrate her philanthropic goals—by other means.

To consider the implications of this issue, let us begin by supposing that there is still only a single kind of investment. In addition to the interest returned to the investor, however, let us suppose that this investment brings others some measure of benefit or harm per unit of resources invested. For each dollar a firm spends today on implementing a method to expand its business, for example, perhaps $0.09 more wealth is generated each year in the future: $0.06 that the firm accrues and returns to the investor as dividends, and $0.03 that accrues to competitors who copy the method.

If a philanthropic investor values wealth in the hands of these competitors as much as wealth spent on his favored causes—as might happen, for instance, if the competitors are among the

intended recipients of his philanthropy—then this circumstance is equivalent to one in which the interest rate is 9% but philanthropists face a disbursement minimum of 3% per year. To a first approximation, if a given philanthropist's ideal spending rate given an externality-free investment opportunity with 9% returns would currently be above 3% per year, then, in the scenario above, he should spend at that rate minus 3%. If his ideal spending rate given such an investment opportunity would currently be 3% or less, he should spend nothing.

Note that in the models of Sections 3.1 and 3.2, a patient philanthropist might do best to spend nothing for some period because of the presence of impatient funders. Nevertheless, we could be assured under those circumstances that the patient party would eventually find it worthwhile to begin spending. Here, we see that when investments carry positive externalities, a philanthropist—in principle one of any time preference, but especially a patient philanthropist—may *never* find disbursement worthwhile, even if there are no less patient funders in the picture.

In practice, of course, a penny "disbursed" to the unchosen beneficiaries of an investment's positive side-effects will typically be much less valuable, from the philanthropic investor's perspective, than a penny he allocates himself. Investments in some firms, however, may produce externalities which are many times larger, in monetary terms, than the returns these investments reap privately. Nordhaus (2004), for example, estimates that innovative firms accrue on average only about 2% of the value they produce. In any event, after multiplying the size of the externality in monetary terms by the extent to which the philanthropist values it less than his own spending, the implications for the optimal disbursement rate are similar.

Finally, just as positive investment externalities should motivate a philanthropist (of any time preference rate) to spend more slowly than he otherwise would, negative investment externalities should motivate him to spend more quickly.

The analysis above assumes that all investment opportunities produce the same externalities. In practice, of course, some investment projects have positive externalities and others have negative externalities. The externalities' magnitudes differ across projects as well.

If a philanthropic investor can choose how she invests her resources without affecting how others invest theirs, she should be willing to accept a portfolio with somewhat lower private returns in exchange for more positive and fewer negative externalities. The practice of seeking out investments with atypically large positive externalities is known as impact investing or ESG (environmental, social, and governance) investing. The practice of foregoing investments with atypically large negative externalities is known as divestment. Both are known as socially responsible investment.

Though these two practices are in principle symmetric, they are typically implemented on the implicit understanding that there is an asymmetry in the distribution of externalities. The value of divestment is typically thought to lie primarily in divesting entirely from harmful sectors, such as the fossil fuel and tobacco industries. The value of impact investing, on the other hand, is typically thought to lie primarily in funding not beneficial industries as a whole but particular firms, typically small and young, offering a new product or service with "outlier" potential to cure some social ill. Thus, whereas standard divestment practices can be implemented by almost any moderately-sized, socially conscious investor, impact investing is typically outsourced to institutions, called *social impact funds*, that specialize in identifying projects with a promise of both private returns and outsized social returns.

Though impact investment and divestment practices are conventionally organized in this way, I am not aware of any explicit defense of the implied asymmetry. Furthermore, as the impact investing industry has grown, it has become more common for social impact funds to include generic investments in seemingly positive-externality industries, such as personal fitness gyms, in the ESG remit. Partly as a result, it is also now common for social impact funds to claim that they offer social benefits at no loss to private returns (see e.g. Geczy et al., 2021). This dilution of the notion of impact investing has drawn criticism from some commentators and industry practitioners, such as Morris (2018), but it might be interpreted as a straightforward extension of the logic of divestment to the case of industries with positive rather than negative externalities.

One investor's behavior does affect others', however. When a philanthropic investor divests from a profitable firm or industry, she opens the investment opportunity to other, less scrupulous investors. Because these investors must now skew their portfolios more toward the given firm or industry than they otherwise would have liked, the philanthropist's divestment does generally decrease total investment in that firm or industry. The size of this decrease is however only a small fraction of the sum by which the philanthropist has divested.[39] Likewise, investing in a profitable firm or industry with positive social benefits largely just crowds out other potential investors.

By the same token, however, the cost to a philanthropist of investing less in socially destructive firms is generally slight. The philanthropist does not forego the available profits altogether; she simply has to skew her own portfolio more toward other firms than otherwise would have been optimal for her. Likewise for the cost of industry-level impact investing in publicly traded firms. Whether the slight benefits are worth these slight costs depends on the circumstances, and cannot be determined by theory alone.[40]

A philanthropic investor can have a more sizable impact on total investment in a given area if the investment opportunity in question is illiquid and idiosyncratic. Consider for example a private equity or venture capital fund—such as a traditional social impact fund—that investigates startups, and decides whether to invest in them partly on the basis of their promise of social benefits. This fund can find socially valuable investment opportunities offering substantially below-market private returns, which would typically not trade publicly in an efficient market. Such opportunities are particularly impactful. If the fund does not invest, the given startup may not raise funding elsewhere, or it may raise less and only after a delay.[41]

Philanthropic investors thus face a natural tradeoff between private and social returns.[42] Taking (or foregoing) investment opportunities with returns at or near the market rate imposes small costs on the philanthropic investor, but the primary effect of such a decision is to crowd out (or in) funding from other investors. A complete analysis of the consequences of these substitutions would be immensely complex. To a first approximation, however, we might

---

[39] See Roth Tran (2019) for a review of the literature on the impacts of divestment in different industries.

[40] See Christiano (2019) for a more thorough exploration of this question. Note that the premise that divestment alters the cost of capital in a given firm or industry is not in tension with the efficient market hypothesis, as several commenters claim.

[41] See Roth (2021) for an analysis of the circumstances under which impact investing is most impactful, relative to return-maximizing investment followed by disbursement.

[42] This point is documented in part by Geczy et al. (2021), and defended by Hillebrandt and Halstead (2018) against optimistic claims to the contrary.

suppose that capital crowded out of (or into) a given industry goes to (or comes from) the rest of the market as a whole. The true externalities reaped by investing in a given project will then lie somewhere between the externalities produced by the project itself and the average externalities of further investment across all industries.

Finally, buying a sufficiently large share of a single firm can give an investor the right to influence what the firm does. (A nonprofit foundation is prohibited from owning more than 20% of the ownership or voting rights of any individual business, but a philanthropist may acquire more than 20% voting rights by not holding them all within a single foundation.) Doing good through "philanthropic activist investing", as this practice might be called, naturally differs in many ways from doing good by shifting the allocation of capital across firms or industries. For our purposes, however, both practices offer the same tradeoff: lower returns (or, through a less diversified portfolio, lower risk-adjusted returns) in exchange for a positive externality (here, the influence exercised by the philanthropist on firm behavior).

By the reasoning outlined in the single-investment case at the beginning of this subsection, this tradeoff, however it appears, is equivalent to a tradeoff between higher (social plus private) returns and higher disbursement minima on the one hand, and lower returns and lower disbursement minima on the other. As explained in Section 3.3, patient philanthropists, more than impatient philanthropists, should generally be willing to sacrifice returns to avoid disbursement minima. Patient philanthropists, therefore, especially those funding public goods primarily funded by impatient parties, should be relatively uninterested in impact investing or divestment, and should more often simply seek the investments with the highest private returns.

That said, recall from Sections 3.1 and 3.2 that the optimal spending strategy in the presence of impatient funders, when investing carries no externalities, is often to invest everything until some point in the future and then to begin disbursing. By the same token, the optimal patient plan in the presence of impatient funders, when at least some investments do carry externalities, may be to seek the investments with the highest private returns until some point in the future, and then to switch to impact investing, either exclusively or alongside disbursement.

4.2 *Mission hedging*

In Section 3 and the beginning of 4.1, we assumed that there was only a single investment opportunity and that it was risk-free. Then, in Section 4.1, we considered how a patient philanthropist should choose among investments offering different social and private returns, but we continued to focus on the case in which these returns were known. We will now explore the implications of risk.

Consider an individual most of whose income comes from freelance work. During periods of high employment, his labor income is high. He consumes some of it and invests the rest. During recessions, he cannot find work, and he consumes by drawing down his investments. When investment returns are known, this person naturally prefers the investment whose private returns are highest. When they are unknown, however, he will gladly accept an investment with lower *expected* returns if these returns are *anti-correlated* with his labor income. That is, he would rather buy stock in a mobile home manufacturer, whose value will be high during a recession (when its business is booming) and low otherwise, than equally-priced stock in a yacht

manufacturer, whose value will follow the opposite pattern. This is the conventional sense in which an investor seeks the investment with the highest "risk-adjusted" returns.[43]

A more general statement of this insight is that, in the face of risk, an asset is more appealing for an investor if its returns are higher when the investor needs money more. For the freelance worker above, money is most valuable during a recession, when he is poorer. For other actors, money may be most valuable under other circumstances. In particular, for a philanthropist, money is most valuable when capital can best be put to philanthropic use.

If the philanthropist's mission is to spend on poverty relief, she should, like the freelance worker, prefer investments that pay off more during recessions, when poverty is severest and each dollar spent on poverty relief goes furthest. Given other missions, too, wealth will typically be at least somewhat more valuable during recessions because most prices will be somewhat lower. As long as a philanthropist is not planning to spend her funds on the immediate relief of social ills, however, the immiseration of the economy in general is probably a poor proxy for the utility of wealth to her mission. Philanthropists, therefore, should typically choose their investments differently from most individuals, who generally face concerns like those of our freelance worker (though perhaps less intensely). The practice of choosing investments that pay off most when wealth is most valuable with respect to a given philanthropic objective is known as mission hedging.

This term was introduced, and the practice was first rigorously explored, by Roth Tran (2019). As Roth Tran observes, there is a common pair of conditions under which mission hedging objectives usually conflict with the impact investing objectives discussed in the previous subsection. These are the conditions that a given industry's production activity causes the problem a philanthropic investor wishes to solve, and that the industry is more profitable when it is more productive. A foundation charged with combating climate change, for instance, might wish not to divest from Exxon Mobil but to invest *more* in Exxon Mobil, on the grounds that scenarios in which Exxon Mobil is profitable are scenarios in which resources for fighting climate change are most needed. A foundation fighting factory farming might do best to invest in Perdue, by the same token, and an antiwar foundation might do best to invest in Raytheon.

By focusing on the (perhaps eyebrow-raising) cases in which the mission hedging and impact investing / divestment considerations point in opposite directions, Roth Tran can compare the considerations' net effects. She argues, I believe successfully, that the mission hedging considerations are probably more important for philanthropic investors in most circumstances. As a result, "mission hedging" has, for some, become nearly synonymous with a recommendation to do whatever turns impact investing or divestment on its head.

Mission hedging considerations, however, can strengthen conventional impact investing and divestment recommendations as well. If an industry is in the business of ameliorating the problem a philanthropic investor also intends to ameliorate, and if the industry is more profitable when it is more productive—i.e. if the first of the two conditions above is reversed but the second is maintained—this industry too will typically be most profitable when the problem is at its worst. An emergency relief fund offering support in the event of epidemics, for instance,

[43] This asset pricing framework, introduced by Breeden (1979), is known as the consumption-based capital asset pricing model. Though it is disputed as an accurate description of typical individual investment behavior, it highlights a consideration of clear relevance to the question of optimal philanthropic investment behavior, as discussed below.

might do well to invest in pharmaceutical manufacturers. Likewise, if an industry causes the problem but is more profitable when it is *less* productive, a mission-hedging philanthropist should divest. In fact, this pair of conditions likely describes the fossil fuel industry. Finding and extracting new oil deposits may be profitable for the firm that possesses them, but an increase to the global oil supply by some proportion lowers the price by a much greater proportion.[44] At least some of the circumstances that bode worst for the climate thus lower, not raise, profits for the industry as a whole. Finally, of course, when *both* conditions are reversed, mission hedging and impact investing / divestment considerations are again in tension.

When a philanthropist aims to do good across a wide and changing array of causes, mission hedging is more difficult to implement. These difficulties are exacerbated when the philanthropist aims to delay his efforts at direct impact until some distant date, by which point the most promising objects of his philanthropy are likely to have changed substantially and unpredictably. There are nevertheless some ways in which, as patient philanthropists, we can profit by mission hedging.

First, as noted in Sections 3.5 and 3.6, we may wish to delay our spending partly in the hope that the resources will be put to better use in rare but nevertheless precise scenarios, such as the advent of a transformative technology or a world war. We can best prepare for such a scenario by searching for investments most likely to have higher returns in the event that the scenario is approaching.

Second, even if we believe that most of our philanthropic resources will best be spent in unforeseeable ways, we can consider broad indicators for long-run trends likely to render the world more or less in need of philanthropy on the whole. As noted above, for example, wealth will typically be somewhat more valuable for almost any philanthropic objective during hard times than during booms. Analogously, when planning on a long timescale, we should probably consider wealth more valuable in the event of long-run stagnation than in the event of long-run growth. This consideration should motivate us to invest similarly to how a typical individual investor invests, at least on the model of the freelance worker described above. We may believe, however, that the inverse relationship between growth and the value of philanthropic resources is stronger or weaker than the inverse relationship between growth and the value of resources for a typical investor. We can then exploit this preference-difference by skewing our portfolios toward the proverbial mobile home manufacturers and away from the yacht manufacturers, or vice-versa.

Likewise, we might believe that the value of our philanthropy will depend on the degree of wealth inequality: a world with more wealth inequality is probably a world with more large and technocratic philanthropists, for instance. Or we might believe that it will depend on the incidence of military conflict. These possibilities are all speculative; I do not believe a serious inquiry has yet been done into the long-run trends likely to render philanthropic resources most valuable. I only mean to illustrate that we do not need to have precise missions, like fighting climate change, factory farming, or risks from artificial intelligence, to successfully engage in mission hedging. We only need to identify a trend likely to change the value of philanthropic resources on the whole *to a different extent* than it changes the value of resources for a typical

[44] See e.g. Bilgin and Ellwanger (2019).

investor, and then to skew our portfolios toward investments whose returns are correlated with this trend, if any can be found.

Finally, and perhaps most importantly, we can monitor the investment behavior of other actors who fund—or, are likely in the future to fund—the projects we intend to fund ourselves. If their portfolios appear misbalanced in some direction, we can then invest so as to shift the collective portfolio as far as possible toward the ideal. Concretely, a substantial fraction of all the wealth earmarked for future-oriented philanthropy is currently invested in a handful of internet technology firms (the wealth of Patrick and John Collison consists primarily of Stripe, for example; the wealth of Dustin Moskovitz and Cari Tuna is largely invested in Meta—formerly Facebook—and Asana) and cryptocurrency trading platforms in particular (e.g. Sam Bankman-Fried's ownership of FTX and Ben Delo's stake in BitMEX). If the cryptocurrency boom collapses, or if any of these firms perform badly, the pool of future-oriented philanthropic resources will shrink substantially, and the value of additional resources in this space will rise.[45] If we can, therefore, we should choose investments whose returns are anti-correlated with those of these firms and industries. Perhaps, for instance, where possible we should simply short-sell these firms.[46] Again, however, as with most mission hedging considerations, this consideration is more action-relevant for philanthropists intending to spend quickly than for those intending to delay. For the latter, a larger share of the total budget that will be allocated along similar lines comes from funding sources that will arrive in the future, and whose investment composition is hard to foresee.

Even if the philanthropic investors whose portfolios we are offsetting have philanthropic goals and relevant beliefs identical to our own, they may be "overinvested" in particular firms (relative to the optimal portfolio for philanthropists with these goals and beliefs) because they face idiosyncratic constraints. In particular, the Collison brothers, Moskovitz, Bankman-Fried, and Delo are co-founders and executives of the firms in which they are so heavily invested. Selling large stakes in these firms would weaken others' trust in the firms' prospects and the executives' incentives. Selling large stakes quickly, furthermore, would incur large capital gains tax liabilities and would temporarily lower the share price. These philanthropists may therefore welcome outside efforts to diversify the collective philanthropic portfolio, and not react by shifting their own investments in the opposite direction. In this case, we can straightforwardly choose the investments that seem best holding the investment behavior of these fellow funders fixed. That is, we can implement the approach I have called BP.

To some extent, however, other funders of the causes we wish to support may also invest as they do because their philanthropic preferences and beliefs do not fully align with our own. Their time preferences may differ, for instance, as explored at length in Section 3. When the investment differences are motivated by genuine differences in objectives, and when at least some of the other funders are sophisticated enough to alter their own portfolios in light of our own as well, we must, as in Section 3.2, recognize that we and the other funders are playing a game. Extending a game-theoretic analysis like that presented in Section 3.2 to the case in which

[45] Case in point: since this document was written, the cofounders of BitMEX, including Delo, have been convicted of violating the Bank Secrecy Act, and Bankman-Fried's FTX has collapsed entirely. Moskovitz's and Tuna's net worths have also fallen substantially along with the values of Meta and Asana.

[46] Further thoughts on how philanthropic investors can find investments uncorrelated, or anticorrelated, with those of other philanthropic investors are offered by Dickens (2021).

funders face multiple investment opportunities, each with its own impact and hedging implications, remains a project for future research.

4.3 *Patient finance*

The previous two subsections focused on the two primary ways in which philanthropic goals introduce investment considerations different from those faced by typical investors. When we considered the interactions between these considerations and time preference, we found that they were generally less significant for patient than for impatient philanthropic investors. That is, patient philanthropists should be less interested in impact investment and divestment than other philanthropists, and they should probably be less interested in mission hedging as well. Instead, their financial behavior should more closely resemble that of patient investors with otherwise conventional objectives.

Patience on its own, however, also should motivate unusual financial behavior. The implications of patience for investment strategy have been explored both empirically—primarily, since Merton (1993), in the context of university endowments, the largest of which are at least somewhat more patient than most investors[47]—and theoretically. An exhaustive review of the literature on patient financial strategy would go beyond the scope of this report, but two considerations appear particularly important and well-established.

First, having a lower rate of time preference should motivate an investor to accept investments with greater volatility in exchange for higher expected returns. Just as it is reasonable for individuals to shift from riskier to less risky assets over their lives, as they approach retirement, "infinitely-lived" and comparatively patient institutions should and do populate their portfolios with volatile assets indefinitely.[48] They can afford to do so, and reap the resulting volatility premium, because a time never comes when the institution will have to spend quickly enough that a temporary drop in its portfolio's value would be disastrous.[49]

Second, patience should lead an investor to make investments that cannot quickly be liquidated. For instance, an investor who may soon need access to his wealth will be less interested in investing in private equity or buying real estate: investments which, due in part to issues of adverse selection, typically cannot quickly be sold without incurring significant investor

---

[47] See Campbell (2011) for a simple introduction to university endowment management. Universities, like philanthropic institutions, are typically nonprofits, which makes their example particularly relevant for tax and legal purposes. Note however that, for the purposes of the previous two subsections, universities are more like typical investors than like philanthropic investors. This is because their goal is not social impact writ large but the prosperity of an individual institution. They have no more incentive to divest from socially destructive firms than the average household, for example—unless of course they are pressured into divestment by students, or other actors with philanthropic objectives that extend beyond the university.

[48] As documented e.g. by Gilbert and Hrdlicka (2015).

[49] Note that the case for volatility-tolerance I am endorsing here is predicated on an investor having unusual patience, not on her having philanthropic goals. Many in the Effective Altruism community, including Tomasik (2015) and Dickens (2020), have further argued that, because philanthropists are typically small contributors to public goods to which many other parties also contribute, they should generally be almost risk-neutral in their investment behavior. While there are conditions under which this advice is sound—such as when donors face idiosyncratic investment opportunities with uncorrelated risks (see e.g. Shulman, 2012)—the correlation of returns across investors implies that philanthropic spending goals alone should typically *not* motivate riskier investment behavior.

losses. Investors who know they will not soon need access to their wealth will thus be rewarded with an "illiquidity premium" for making these investments.[50]

Partly as a result of these investment strategies, the five largest US university endowments dramatically outperformed US and UK equity returns over the fifteen- and twenty- year periods preceding June 2016 (Azlen and Zermati, 2017). By contrast, they underperformed US equities from 1900 or 1950 to 2017 (Chambers et al., 2020), due largely to worse performance before the 1980s, when they first began to act on the lessons outlined above (Campbell, 2011).

The story is admittedly more complex than these facts might suggest. If we consider all endowments holding more than $1b, and consider only the 15 years preceding June 2016, Azlen and Zermati find that they slightly underperformed US (but not UK) equities. Moreover, university endowments have performed relatively badly in the past few years (Flood, 2019). In any case, large university endowments have certainly outperformed the global average portfolio, which consists roughly of 60% equities and 40% bonds, over any reasonable horizon.

4.4 *Movement building*

As we have seen repeatedly, a central difference between philanthropists and most other investors is that philanthropists fund projects to which other parties may also be interested in contributing. This observation suggests an "investment opportunity" uniquely available to philanthropists: fundraising.

In the simplest case, spending a dollar on fundraising can produce a stream of future increases in donations by others to one's own philanthropic institution, or to projects the institution otherwise would have funded. The rate of return to fundraising is the interest rate this dollar would have had to earn in order to produce an equal stream of inflows. To a first approximation, the institution should spend on fundraising activities instead of engaging in conventional financial investments as long as the rate of return from the former exceeds the rate of return from the latter.

A more complete comparison of fundraising and financial investment would have to account for the fact that fundraising only directs wealth from one use to another, whereas other investments typically produce new wealth. This complication is exacerbated by the fact that, if a funder is willing to contribute some funds to one's cause, she would likely have spent those funds in a relatively valuable way (from one's own perspective) even in the absence of any fundraising effort. If one targets one's fundraising efforts carefully, though, or believes that spending on one's chosen projects is many times more valuable than spending in adjacent domains, the returns from fundraising and other investment activities can be compared straightforwardly.

To take this reasoning a step further, a philanthropist can also direct others' resources toward his own objectives by persuading others to join the community of individuals who are passionate about these objectives and contribute to them on an ongoing basis. Religious institutions, for example, in addition to typically being intrinsically interested in gaining new

[50] Dickens (2022) agrees that these are the two primary ways in which patience should motivate unusual financial behavior, but argues that the benefits of acting on these unusual liquidity and volatility preferences are typically too small to matter in practice.

members, gain financially from the tithes and labor services that new members often contribute. In part, therefore, expenses on proselytization function as investments. Likewise, some philanthropists in the Effective Altruism community spend on efforts to encourage others to take the Giving What We Can pledge, a commitment to give at least 10% of one's income on Effective Altruist principles.

Unless movement members contribute their resources directly to the institution, however, or allocate them precisely as the institution most prefers, the resources moved through movement-building efforts are not as valuable as resources acquired through investment, and must be discounted accordingly. This discrepancy is likely to be particularly acute for patient institutions. If an institution considers it its optimal plan to invest for centuries before disbursing, and if movement recruits do not do the same with their own wealth, movement building is analogous to financial investment only if the recruits themselves put all their efforts into growing the community further for many generations. If some fraction of recruits' efforts will instead be spent on projects with immediate impacts, then movement building is, like impact investing, analogous to an investment that may offer higher returns but is subject to a disbursement requirement.

That said, there is another reason why a community of philanthropists might want to increase its membership. This is that some of the tasks necessary for the accomplishment of its goals may require the labor of ideologically sympathetic individuals, and not be achievable with capital alone. For instance, though a philanthropist can make grants to others doing work she considers valuable, even if these others don't entirely share her values, she cannot easily delegate grantmaking itself to those with whose priorities she disagrees. As she and her community grow their capital over time, therefore, they must also maintain or grow the community of people genuinely interested in spending this capital in similar ways.

Trammell and Sempere (in progress) study the implications of these two motivations for movement building—i.e. the fundraising and specialized-labor-production motivations—in more detail. Under certain assumptions, we solve for a patient social movement's optimal growth path of spending and labor effort on "direct work" and "recruitment". One natural assumption is that unused capital grows at a positive rate whereas, in the absence of active efforts at recruitment and intergenerational value transmission, movement membership shrinks or grows more slowly. As a result, because of the complementarity between labor and capital, a patient movement on the optimal path will, at least in the long run, encourage its members not to take high-paying careers and contribute a share of their incomes but to engage entirely in direct work or recruitment. The movement's expenditures will be funded entirely from a proportionally slow-spending, and therefore fast-growing, endowment. We also find that the optimal long-run disbursement rate may be higher or lower in this setting than in a model in which good is achieved with capital alone, but that it is probably similar.

Of course, a philanthropic community or social movement may have many other motivations to expand its membership. It may hope to use its larger numbers to wield political influence or change social norms, for instance, or it may believe that recruitment saves the recruits' souls. Here, we have focused only on the motivations that appear most directly relevant to the disbursement scheduling question. Broader discussions of social movement theory and practice are given by McAdam and Snow (2009) and della Porta and Diani (2016).

4.5 *Implications for disbursement scheduling*

We noted in Section 4.1 that positive investment externalities should motivate a philanthropist to spend more slowly, and negative investment externalities should motivate him to spend more quickly. We also noted in Section 4.4 that introducing a more complex model, in which capital and "specialized labor" are both necessary for social impact, could motivate faster or slower spending in the long run than is optimal in a model in which only capital is necessary.

In sections 4.2 and 4.3, and the first half of 4.4, we explored ways in which patient philanthropists—because of their patience or their philanthropy—have access to investments that offer higher returns, in some sense, than the investments other parties are able or willing to make. We did not, however, discuss whether the availability of these investments should motivate faster or slower spending than would be ideal if patient philanthropists had to hold their wealth in portfolios like those held by more typical investors.

It may seem at first as though the presence of superior investment opportunities makes investment more attractive relative to spending. In fact, though, there are two effects at play. Higher interest rates let us direct more resources to our chosen causes in the future at a smaller sacrifice to the present. But they also imply that we will have more resources in the future, and thus that there is less value in setting aside additional resources.

If the impact of increasing spending on our chosen causes diminishes slowly as our absolute spending rate grows large, better investment opportunities should motivate us to spend more slowly. If the impact diminishes quickly, better investment opportunities should motivate us to spend more quickly. If our invested resources will ultimately be spent on individual consumption, or on a cause for which the shape of the relationship between spending and impact is similar to the apparent relationship between individual consumption and welfare, then, for these purposes, the impact of increasing our spending in the future "diminishes quickly". On balance, that is, the existence of the patient philanthropic investment opportunities discussed throughout most of this section should probably lead a patient philanthropist to spend somewhat more quickly than would seem ideal otherwise—or, to begin spending sooner than would seem ideal otherwise.

Either way, though, they should not lead a patient philanthropist to spend so quickly that her wealth grows as slowly as that of impatient philanthropists and more typical investors.

5. Patience in history

When we invest for the sake of future philanthropic spending, we run the risk that our resources will not be put to good use. Our successors, or even our future selves, might spend in ways we would disapprove of today. The resources might be expropriated by a future government or destroyed in a future war.

If these contingencies are equally likely to befall patient philanthropists and more typical investors in any given period, they should not change any of the reasoning of Sections 3 and 4. When investors in general face a 0.5% risk of expropriation between one year and the next, this will increase every party's (non-pure) time preference rate that year by 0.5%. Interest rates will

have to be 0.5% higher to render impatient parties indifferent between investing and borrowing, and this will compensate patient parties no less.

Indeed, recall the long-term spending plans proposed in Section 3.3. They are predicated on the assumption of a time preference gap, between patient philanthropists and others, of 1.5% per year. More precisely, they are predicated on the assumption that interest rates reflect time preference rates of 2% per year, and that 0.5% of these consist not of pure time preference or mortality risk but of factors relevant to patient philanthropists as well, like those listed above. (The 0.5% figure is taken from analyses presented in the following subsection.) That is, I have assumed that the life expectancy of a patient philanthropic investment vehicle would be about 200 years. The recommendation to invest for centuries is thus reached despite the acknowledgment that an attempt to do so would probably fail—and despite the assumption that, in the event of failure, funds would be put not merely to less valuable uses but to no use of value at all.

Slow-spending philanthropists, however, may face challenges over the long run that others do not, due to their slow spending or to their philanthropy. They may also face unique opportunities. These considerations could indeed change the recommendations of earlier sections. To evaluate them, we must study how patient philanthropic behavior, and patient financial behavior more generally, has played out in the past. In light of the lessons of history, and with the aid of careful reasoning when the historical lessons are ambiguous, we can begin to form more complete judgments of how patience will unfold in the future.

5.1 *Institutional mortality*

The most straightforward way in which philanthropists' long-term investment plans can fail is that their institutions can collapse or be expropriated. All institutions, and even households, face this risk to some extent. Nevertheless, because it may be higher or lower for philanthropic institutions than for other actors, it is necessary to estimate the risk rate for philanthropists explicitly. If philanthropic institutions face higher "mortality risks" than other actors, and if this gap in the risk rate exceeds the extent to which pure time preference and mortality risk motivates most people to act impatiently, then the reasoning of Sections 3 and 4 is largely reversed. In this case, the risk of losing all one's resources in any given year outweighs the potential benefits of having delayed. Patient philanthropists should still spend more slowly (or later) than philanthropists who do have pure time preference or a desire to spend within their lifetimes, but they should spend *more* quickly (or sooner) than non-philanthropic funders.

A related issue is that the mortality risk of a philanthropic institution may be *age-dependent*. Most importantly, it may increase with age. If it does, we might expect that philanthropists who recognize this tendency can escape it by stipulating that their wealth periodically be transferred to younger foundations with similar goals. For one reason or another, however, this may be infeasible: for example, the future may produce no young foundations that would hew closely enough to the original founder's vision. In this case, even if patient philanthropists would do

best to delay spending in the short term, plans for very-long-term philanthropic investment may not be optimal.

The modern, "American-style" foundation is too recent a development for there to be much evidence on the level or age-dependence of its mortality risk. Almost all were founded within the last century, and there are no examples of large collapses or expropriations, at least within the United States.[51] Nevertheless, we may be able to shed light on the mortality risk of philanthropic institutions by studying the historical longevity distributions of other categories of institutions.[52]

Given a dataset of institution start and end dates, we can run a *Weibull regression* to estimate (a) the "life expectancy" of institutions in this class and (b) whether they appear to close at a constant rate, a rate that grows with time (aging), or a rate that falls with time (child mortality). We can also study whether a given institution-class has been growing shorter-lived over time, longer-lived, or neither.

To begin with the largest-scale institutions on record, Taagepera's (1978a, 1978b, 1979) dataset of empires through the past five millennia suggests a life expectancy of 348 years. Furthermore, the longevities are distributed essentially exponentially, as they would be if they exhibited neither child mortality nor aging. (See Figure 10 below.) Finally, they do not appear to exhibit a trend of shortening or lengthening over time. Empires, at least on Taagepera's definition, can thus be modeled as exhibiting a constant mortality rate of 0.29% per year.

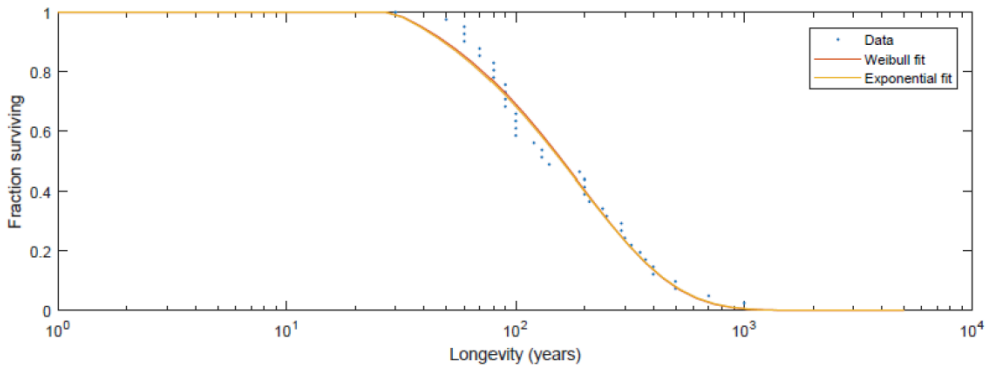

Figure 12

A similar though less formal analysis by Freeman (2020) produces an even longer empire life expectancy estimate of 382 years. His numbers imply that empires do exhibit aging, but the

[51] At least of foundations established before 2009; see Flishman, Kohler, and Schindler (2009). As documented by Levine et al. (in progress), only four of the foundations in their dataset had closed as of their writing or have closed since. Three of these are the Rosenwald Fund (1917-1948), Olin Foundation (1953-2005), and Atlantic Philanthropies (1982-2020), all of which closed for reasons of donor intent. (Recall from Section 1.1 that Atlantic Philanthropies was run by Chuck Feeney, coiner and promoter of "giving while living".) The fourth is the General Education Board (1902-1964), which was subsumed into the Rockefeller Foundation.

[52] Except where otherwise stated, the following analyses were collected or performed by Sandberg (in progress).

effect is mild and nearly nonsignificant. Wikipedia's list of empires implies a life expectancy of 271 years, without child mortality or aging, or 0.37% annual mortality.

Looking on a somewhat smaller scale, a Schönholzer and Weese (2019) dataset of European states from 1000 to 1850 implies a shorter life expectancy of 137 years. This data exhibits moderate child mortality, however, implying a long-run mortality rate lower than the 1/137=0.73% that would be implied by constant mortality.

Smaller still, and perhaps more relevantly, Levine et al. (in progress) conduct this analysis for a dataset of European universities. We find a life expectancy of 361 years and nearly nonsignificant child mortality, implying an annual mortality rate of 0.28%. We likewise analyze the longevities of Catholic religious orders, finding an extraordinary life expectancy of 1,055 years.[53] This data exhibits substantial aging, however. Most of this apparent aging is likely due not to aging *per se* but to the long-term decline in vocations to Catholic religious life, but the lesson is similar. When founding an institution intended to endure, one must consider the possibility that the institution itself, or the class of institutions to which it belongs, will grow harder to sustain as the institution ages.

Probably most relevant for our purposes would be a longevity analysis of *waqfs*, or Islamic trusts. As documented for instance by Kuran (2001), the waqf is often identified as the primary precursor to the later European trust and the modern foundation. These charitable institutions provided public goods and poverty relief throughout the Muslim world for many centuries, and were essentially all intended to be permanent. They grew, both in number and in wealth, until they commanded remarkable proportions of the resources of the societies in which they were established. In the late 1700s, about 20,000 waqfs operated in the Ottoman Empire, collectively earning an annual income equaling one third the Empire's tax revenue. In the mid-1800s, waqfs owned one third of the arable land in Tunisia and well more in Algeria and Greece. In 1923, they collectively owned one eighth of the arable land in Egypt, one seventh of that in Iran, and an astounding three quarters of that in Turkey.

It is sometimes argued that the history of waqfs should discourage us about the prospects for attempting long-term philanthropic investment today. Indeed, though some relatively small waqfs endure to this day, the importance of the waqf system has greatly diminished, due largely to several waves of expropriations over the last two centuries, such as one following (and partially during) the collapse of the Ottoman Empire. Without a longevity analysis like those described above, however, it is far from clear that waqfs were especially vulnerable to expropriation or exhibited aging. The broad strokes of the history are also compatible with the conclusion that waqfs have merely exhibited rare and age-independent, but highly correlated, expropriations.

One relevant data source is Çizakça's (1993) collection of records on *cash waqfs* in Bursa from 1555 to 1823. These were waqfs endowed with cash rather than land, as was more common.

[53] Their unusual longevity has motivated various qualitative explorations of the implications of monastic governance for corporate governance: see Rost et al. (2010), Inauen et al. (2010), and Mercier and Deslandes (2017).

As Kuran (2001) notes, cash waqfs had to survive by lending at interest, which was not generally permitted under Islamic law. They were thus subject to punishing regulations that other waqfs were not, and this greatly weakened their long-term viability. Nevertheless, Çizakça (2004) reports that 20% of those in his 1993 data lasted over 100 years. If their mortality rates were constant, this implies that they closed at a rate of 1.6% per year: still lower than standard estimates of typical time preference rates, albeit only slightly so. In any event, a proper longevity analysis of the Bursa cash waqf data would be valuable and feasible, as would analyses of historical waqf data more broadly.[54]

Firms constitute a final class of institutions whose longevities have been studied. They are far shorter-lived than any of the institution-types above: Daepp et al. (2015) find that modern corporations exhibit constant mortality rates of 10% per year, lasting an average of 10 years.

Unlike most states and nonprofit institutions, however, most firms are not in any sense intended to endure indefinitely. When a firm's business model is no longer as profitable as an alternative use of its capital would be, its owners have every reason to sell it or shut it down. The rate of firm turnover therefore tells us little about the viability of long-term philanthropic investment.

Furthermore, Daepp et al. find a long "tail" of unusual firms exhibiting lifespans averaging 170 years, or annual mortality rates of 0.59%. Rose (2020) discusses these as well, finding that they are disproportionately concentrated in brewing or hospitality and disproportionately German or Japanese. This suggests that, at least in industries and environments with a special concern for maintaining tradition, firms can be much longer-lasting.

In sum, though more research on institutional longevity would be valuable,[55] my present judgment is that philanthropic institutions intended to endure are reasonably well modeled as facing a roughly age-independent long-run closure rate on the order of 0.5% per year.[56]

5.2 *Institutional quality*

[54] The only other waqf dataset of which I am aware is Adiguzel and Kuran's (2021) dataset of Istanbul waqfs from 1457 to 1923, which does not appear to be public. Many waqf records throughout the Middle East and North Africa remain undigitized and unstudied.

[55] One alternative approach would be to attempt to extrapolate from the global history of expropriation risk estimates produced by Political Risk Services (PRS) since 1993. PRS is a consulting firm that advises firms on the political risks to doing business by region, such as the risk of full expropriation or the risk of forced contract renegotiation. This data is widely used for academic research on institutions. Acemoglu et al. (2001), for example, used this data to indicate "quality of institutions today" when attributing bad present-day institutions in former colonies to extractive behavior by the colonists.

[56] This foundation closure rate would be the median of the historical closure rates of Catholic orders, European universities, empires, foundations, European states, long-lived firms, and Bursa cash waqfs. Note however that many of these institution-classes have exhibited low mortality despite a relative lack of explicit research into the determinants of institutional longevity, as illustrated for instance by the lack of investment in quantitative analysis of waqfs. We might expect that an institution established following such research would be at least somewhat longer-lasting.

Even if a philanthropist's institutions endure, his plans for them may not. His successors may not be as effective at grantmaking as he is, or he himself may grow less effective with age. Likewise, his values in the future, or the values of his successors, may differ from his current values in ways that lead them to make worse funding decisions, from his current perspective. Both developments might be called declines in "institutional quality", and both should motivate a philanthropic institution to spend more quickly.

The risk of a decline in institutional quality is in some ways symmetric to the possibility that, if a philanthropic institution closes before disbursing all its capital, the capital it has accumulated will still be put to some good use. As noted at the beginning of this section, if we take the philanthropic time preference rate to equal the mortality rate for philanthropic institutions, we are implicitly assuming that an institution's resources will be spent fully effectively as long as the institution lasts, and do no good at all after passing into other hands.[57] In practice, the difference in value may be less discrete. Accumulated resources that are, say, transferred to public coffers in a sharp tax hike may not be entirely wasted, from the founding philanthropist's perspective; and they may eventually fail to be used well even if they remain within the institution.

Alternatively, a philanthropist may believe that her institution's "quality" will improve over time. She or her successors may grow more effective or enlightened than she is today. (Presumably she will not expect expropriated resources to do more good than resources in her own or her successors' hands, or else she would plan for resources to be transferred to the expropriating party.) In this case she should spend more slowly than would seem best otherwise.

As usual, what matters for the qualitative argument of earlier sections is whether, in the long run, considerations of declining institutional quality are more severe for philanthropists than for other investors, and if so, whether this outweighs the time preference differences discussed in Section 2.

As noted in Section 1.1, some argue that young foundations spend more effectively than old ones. If this is true, we might expect that philanthropists who recognize this tendency can escape it by stipulating that their wealth periodically be transferred to younger foundations. If for some reason this is infeasible, however, and if the relative ineffectiveness of old foundations is due to age itself rather than to improvements in the quality of new foundations over time, then foundations—and perhaps philanthropic institutions more generally—should spend more quickly than would otherwise seem best. In particular, if a foundation's effectiveness is expected to fall at a constant rate each period relative to that of other financial market participants, foundations should simply spend as if they employed a correspondingly higher time preference rate. If the resulting rate is higher than the time preference rate of truly impatient actors, the analyses of Sections 2–4 should be reversed.

I am not aware of an attempt to quantify the effectiveness of foundation spending over time. Such research would be valuable. In the absence of empirical evidence, however, the assumption of a constant "depreciation rate" in institutional effectiveness—let alone one of say, 1.5% per

[57] In the language of Section 2.1, we are assuming that philanthropic institutions have no "bequest motive".

year—strikes me as extremely pessimistic. It would essentially imply that, all else equal, an institution founded by some half-remembered philanthropist 150 years ago would be able to recruit grantmaking talent five times better than one founded 250 years ago, and so on through the centuries. If a long-lived foundation does indeed stay active and continue to hew to its original values, even if it loses its initial dynamism, I would argue that it is most reasonable to assume that the effectiveness of its spending typically falls not to zero, as it would if it fell at a constant rate, but to a positive fraction of the effectiveness in early years. In this case, giving in early years offers only a temporary "multiplier", like that offered by supporting the Ugandan school of Section 2.6, which the time preference gap eventually overwhelms. The conclusions of Sections 2–4 then hold unchanged.

Furthermore, this consideration must be offset by the extent to which, as time goes on, we may continue to develop more effective ways of spending money to do good. As noted in Section 3.5, many of the most important funding areas today were unavailable or inconceivable to funders in the past. Local almsgiving has always been an option, but today philanthropists also regularly send resources to the poorest parts of the world, fund pathbreaking scientific research, and use randomized controlled trials or careful theoretical analyses to improve their interventions. In aggregate, these developments may be thought of as representing a dramatic improvement in the quality of philanthropic institutions. The future may hold further such improvements.[58]

Other investors seem unlikely to enjoy analogous prospects. Machine learning systems may offer households ever better purchase recommendations, say, but to date such innovations have presumably not improved the effectiveness with which households spend their consumption budgets nearly as much as social scientific research has improved the effectiveness with which we can spend our philanthropic budgets. This is not surprising; the more circumscribed one's goals, the likelier it is that they will be achievable without decades or centuries of research. The space of ways to spend on oneself, given one's tastes and circumstances, is typically well understood. The space of all possible ways to spend on doing good is still far from fully explored.

Though I am skeptical that the quality of a philanthropic institution should be expected to degrade indefinitely in the long run (conditional on its survival), there are good reasons to suspect that philanthropic institutions, especially when acting patiently, face unusually severe risks of long-run "value drift".

The most straightforward source of value drift risk is simple reversion to the mean. Most people care, or spend as if they care, far more about themselves and their loved ones than about the world at large: even in the United States, by far the most charitable developed country,

[58] As of this writing, users of the forecasting platform Metaculus predict that the most cost-effective way to save a life, or do a similar amount of good, identified by charity evaluator GiveWell in 2031 will cost substantially less than the most cost-effective way to do so identified by GiveWell today (Metaculus, 2021). Likewise, a recent survey (Hoeijmakers, 2020) of nonprofit leadership in the Effective Altruism community finds that most believe that, all things considered, their institutions will spend resources substantially more effectively in 10 years than today, due primarily to further insights along these lines.

charitable giving constitutes only about 2% of GDP.[59] There is thus a clear risk that one will eventually abandon one's commitment to philanthropy and spend one's wealth self-servingly, or that one's successors will do the same. Other investors do not entirely escape this concern; for instance, individuals face some risk that they or their descendants will gamble their savings away. But it seems safe to say that people can more reliably be trusted to spend money well on themselves than to stay true to a philanthropic plan.

Even if a philanthropist and his successors remain committed to spending philanthropically, there is a risk that resources will be spent in pursuit of philanthropic values at odds with the founder's. Just as value drift appears particularly concerning for philanthropic institutions in general because of the relative rareness of philanthropic motivation, philanthropists with particularly idiosyncratic values—or rather, with values that they believe will be idiosyncratic in the future—ought to be particularly concerned that their resources will eventually be put to uses that do not further those values.

A particularly relevant source of value drift risk for patient philanthropists is the risk that the philanthropist or her successors will grow less patient. If the benefits of delaying spending grow monotonically until the patient-optimal spending time, this risk can shrink the expected benefits of delaying spending, but it cannot make them negative. That is, if it is best to begin spending a certain sum in 200 years, somewhat worse to begin spending in 100 years, and worst of all to begin spending today, then whatever the risk that successors will begin spending too early, it is better today to delay than to spend. However, the expected benefits of delaying spending may not always grow monotonically. In the Ugandan school scenario of Section 2.6, for example, and in the discussion above about the hypothesized dynamism of young philanthropic institutions, the case for waiting relies on a long-term trend outweighing a short-term multiplier. In these cases, even though beginning to spend in 200 years is better than beginning today, beginning in 100 years may well be worse.

Trusts can mitigate value drift risks by adopting narrow missions with precise beneficiary groups. These beneficiaries then have the power and motivation to sue the trust if its resources are not spent for their benefit, just as shareholders can and are incentivized to discipline a misbehaving CEO. This remedy is almost entirely unavailable for an institution designed to be flexible enough to address the problems of the future, to spend more slowly than its beneficiaries would like, or to wait a long period before disbursing at all. Such institutions must, like most foundations, essentially remain accountable to no one but themselves, and rely on a faithful line of successors to execute the founder's intention.

If there is a certain probability each period that a philanthropic institution will switch from fully acting upon its founder's values to pursuing goals the founder would have considered worthless (such as luxurious lifestyles for its own executives), then value drift risk, like expropriation risk, can be modeled as a simple increase to the philanthropist's time preference rate. In general, however, value drift risk introduces complexities that cannot be modeled this way.

[59] Giving USA (2021); an open-access summary of statistics up to 2017 is given by Charity Navigator (2018).

Consider the following stylized example, adapted from Negele (2020). A founder wants to advocate for the income tax rate on the wealthy that is optimal given equal concern for the welfare of the rich and the poor. Resources taxed from the rich will be transferred to the poor, but taxes on the rich will disincentivize wealth creation. The economy is stagnant from generation to generation.

Because the research is so ambiguous, the founder expects his guess about the welfare-maximizing tax rate to be off by about 8% in one direction or the other. The loss from getting the rate wrong is *quadratic*: being off by 2% is 4x as bad as being off by 1%, being off by 8% is 64x as bad, and so on. If the founder delegates the decision to an expert in the next generation, he believes that this successor will have access to better research: her guess will be off by only about 5% either way. He also believes, however, that this expert will care more for the rich than the poor, such that the optimal tax rate from her perspective will be 5% lower than the optimum from his own perspective. By delegating to her, there is a possibility that her "bias" and her "error" cancel out, so that the rate she chooses is the optimal rate on his view; and there is an equal possibility that the rate she chooses is 10% lower than the optimal rate on his view. Since $10^2/2 = 50$, which is less than 64, delegation seems like the right decision.

Suppose however that this successor also has the power to delegate to a successor. This second successor will be perfectly informed, but he will favor the rich even more than the first: his optimal tax rate will be 4% below hers. From the first successor's perspective, acting herself incurs an expected loss 25x greater than that of getting the tax rate "wrong" by 1%, whereas delegating incurs a certain loss only 16x greater. She thus re-delegates. The founder's resources ultimately lobby for a tax rate that is 5% + 4% = 9% lower than the rate the founder would have considered optimal, incurring a loss greater than what he would have incurred by acting immediately (since $9^2=81 > 64$).

Under some circumstances, therefore, the dynamics of re-delegation can make delegation more dangerous than it would otherwise seem. Under other circumstances, by the same token, they can make delegation safer: it may be worthwhile to delegate to a successor, even if one would not want her to make an important spending decision herself, because one knows that she is likely to redelegate to someone more benign. The lesson is simply that the value drift risk incurred by long-term philanthropic investment plans can differ from a simple accumulation of those incurred by a sequence of short-term investment plans. This phenomenon is explored more broadly in the growing economic literature on "policy drift" (which we may treat as roughly synonymous): see Gieczewski (2021) for a recent example and review of the literature. It would be valuable to consider the theoretical implications of drift for the design of philanthropic institutions more thoroughly than has yet been done.

Value drift is difficult to quantify, and I am not aware of any systematic attempts to measure it at the institutional level. Average individual attrition rates have been calculated for religions and social movements in various contexts,[60] but institutions often slow these trends by filling their

[60] Attrition in the Effective Altruism community may be particularly relevant to our purposes. An informal survey by Savoie (2018) suggests that roughly 10% of those highly involved in the Effective Altruism

leadership roles only with individuals displaying strong evidence of long-term commitment to the institution's values and by making many hiring, dismissal, and policy change decisions by committee rather than unilaterally. Consider the Catholic Church, for instance, whose commitment to a given point of doctrine is plainly much longer-lived than the average Catholic's commitment even to the basics of Catholicism. Holcombe (2000) likewise argues that, though American foundations generally exhibited a drift toward the political left during the 20th century, this trend did not encompass foundations that identified themselves explicitly as conservative and hired with a view to preserving this identity.

Indeed, the consensus among historians of philanthropy appears to be that philanthropic institutions have historically adhered to narrow missions *too* faithfully as the needs of the world have changed. There are countless examples of age-old foundations and trusts successfully constrained by the "dead hands" of their founders to support niche, and eventually (arguably) "overfunded", causes. The Haseki Sultan waqf—established in 1552 to support one soup kitchen, one mosque, and two hostels in Jerusalem—came to possess 26 entire villages and numerous shops, mills, and bathhouses. ScotsCare, established in 1611 as the Royal Scottish Corporation for the benefit of poor Scots who moved to London following the union of Scotland and England, enjoys a £52m endowment[61] still reserved for the benefit of Scots in London, though Scots are undoubtedly now one of the city's wealthiest ethnic groups. Holcombe (2000) thus presents the all-purpose American foundation—an institution inaugurated in 1913 with Rockefeller's mission to "promote the well-being of humanity throughout the world"—as an improvement on the model of the more narrowly targeted European trust. Kuran (2001), in turn, goes so far as to suggest that the rigidity of waqfs, relative even to European trusts, contributed to the relative decline of the Islamic world.

To avoid this trap, a philanthropist might do better to define his institution's mission in terms of values and principles rather than concrete projects. Even so, though, he may err by defining those principles too strictly to allow for moral progress. The Rhodes Trust, for instance, was originally intended to support only male scholars, but it no longer discriminates on the basis of sex. Its resources are thus presumably better spent now—or at least spent according to better values—than they would have been if Rhodes and his successors had spent down the trust more quickly.

In the best-case scenario, therefore, waiting to spend resources secures an improvement, rather than merely no degradation, in the values by which the resources will be spent. Just as the effectiveness of philanthropic institutions may improve with time, giving them reasons to delay spending not shared by most other investors, the values of philanthropic institutions may improve with time as well.

---

community leave it each year. Todd (2020) presents retention data from several particularly engaged subgroups of the EA community, finding attrition rates lower than Savoie's but still higher than a conventional time preference rate.

[61] https://www.peridotpartners.co.uk/wp-content/uploads/2021/01/Our-2019-Accounts.pdf

This is not to say that, for philanthropists, more flexibility is always better. The point is just that, when a founder wishes to bind his resources to particular philanthropic uses, it appears he can typically do so, as long as the legal regime surrounding his institution lasts. Patient philanthropists thus face a real risk of tying their heirs' hands too tightly, and a real tradeoff between excessive drift risk and excessive mission rigidity. One of the Rockefeller Foundation's first projects was the eradication of hookworm in the American South, largely achieved from 1910-14.[62] It is a good thing that the Foundation was not devoted exclusively to hookworm eradication in the 1960s: such a restriction would have precluded the Foundation-funded "Green Revolution", ramped up at the time, which is estimated since to have saved over one billion lives.[63] In 2016, by contrast, when hookworm resulted in a loss of approximately 2-4 billion disability-adjusted life-years worldwide[64] and the Foundation spent $6m sending children to live performances of Hamilton,[65] the sign of the value of this flexibility was less clear.

In my judgment, absent more rigorous evidence on the trajectories of philanthropic institutions in various settings, history suggests that the balance is best struck by adhering to a precise but broad set of principles, preferably associated with a stable community of enthusiastic adherents. Religious institutions seem clearly to benefit from the endurance of an ideologically sympathetic community, for instance, which can refine the expression of its values across the generations without abandoning its more fundamental principles, and from which the institution's managers can be drawn. Institutions predicated on environmentalist or libertarian principles are endowed with similar advantages. The Effective Altruism community, or parts of it, may come to serve a similar role for institutions committed to pursuing the classical utilitarian objective. Even if one finds that adopting such one-dimensional religious, political, or philosophical commitments cannot fully capture the nuances of one's own philanthropic goals, this price may be worth paying in exchange for the ability to invest capital relatively faithfully over the relatively long run. As discussed in Section 4.4, sustaining and growing such a community from period to period may in turn be a highly valuable activity even for a philanthropic institution intending to delay most of its impact.

All things considered, there does not yet seem to be a clear answer to whether the "quality" of a philanthropic institution, designed as carefully as possible to flourish over the long run, should be expected to improve or degrade over the centuries after its founding.

5.3 *Institutional growth*

In Sections 5.1 and 5.2, we explored the possible dependence of an institution's closure risk or quality on its age. In general, however, a long-lived philanthropic institution should aim not only to endure but to grow. Indeed, as we have seen, much of the value in establishing an institution

[62] https://rockfound.rockarch.org/eradicating-hookworm
[63] Center for Strategic Philanthropy and Civil Society (2007)
[64] Bartsch et al. (2016)
[65] Rockefeller Foundation (2016)

that invests for long periods comes from the ability of its asset growth to outweigh the considerations that favor spending more quickly. If the longevity or quality of an institution increase or decrease in its size, the case for long-term philanthropic investment can therefore be strengthened or considerably weakened.

Unfortunately, I am aware of even less formal research on size-dependence than on age-dependence in these areas. Until such research is done (or comes to my attention), therefore, the discussion of this subsection must, like that of the previous subsection, rely largely on informal arguments and anecdotal observations.

A first observation is that small nonprofits, like small firms, typically find it more difficult to recruit and retain top management, grantmaking, and investing talent than larger institutions, even after accounting for the fact the latter can generally afford better compensation. This is because philanthropically-motivated individuals are "compensated" in part by the prospect of having a large impact with their careers, and because individuals in all industries are compensated in part by the experience and prestige that comes with size. Large investors also have access to investment opportunities that small investors do not, so a large philanthropic investor can better take advantage of some of the idiosyncratic investment advice offered in Section 4.

On the other hand, of course—and for the same reasons—a larger pool of resources is more likely to attract opportunists aiming to divert these resources to their own ends. Likewise, it is more likely to motivate expropriation by the state.

That said, a larger fund is also better able to afford the internal infrastructure, political influence, and so on needed to lower these risks. Consider that larger corporations typically pay lower tax rates than smaller corporations, at least in the US (Hager and Baines, 2020). Also, as reported in Section 5.1, empires seem to have exhibited substantially longer lifespans than (European) states as a whole. A plausible interpretation of this pattern is that, at least in the context of corporate or governmental institutions, size gives one the resources to protect oneself from "threats", and this benefit usually outweighs the extent to which size motivates more attempts at infiltration or expropriation (including taxation).

A complete account of the relationship between the size of philanthropic institutions and their longevity or quality would have to incorporate many further considerations. Larger foundations, just like larger firms, can for instance benefit from economies of scale; smaller ones can for instance benefit from an absence of bureaucracy. In light of all these considerations, the relationship between the size and the effectiveness of a philanthropic institution pursuing a given class of goals, all this considered, may well not be monotonic, just as the relationship between firm size and firm profitability in most industries is not monotonic. That is, it is likely possible for such an institution to be too small or too large. If so, these institutions achieve their goals most effectively by splitting as they grow, so that they are each always approximately the optimal size. Being too small is thus a reason to spend more slowly and grow more quickly, but being too large is not a reason to do the reverse.

This observation may also make long-term investment plans appear more realistic. The idea that a single philanthropic institution could endure and grow for centuries until it owns a substantial fraction of the world's assets may seem bizarre. As the millennium-long history of the waqf system in the Islamic world demonstrates, however, a network of smaller philanthropic institutions can indeed, under the right circumstances, endure and grow to similar proportions.

Some of the political risks and costs associated with the growth of philanthropic institutions, however, may depend not on the growth of the individual institution but on the growth of the industry as a whole. These cannot be mitigated by managing individual institutions' sizes.

The most straightforward political, industry-wide obstacle to the investment-based growth of philanthropic institutions in the United States is the foundation disbursement minimum. Since the Tax Reform Act of 1969, American foundations have enjoyed tax privileges only on the condition that they disburse at least 5% of their assets per year. If they face an interest rate of 7%, this implies a maximum growth rate of 2%, which is roughly the economic growth rate. The disbursement requirement thus essentially requires foundations to spend as "impatiently"—and grow as slowly—as the average, typically impatient, actor. As noted in Section 3.3, this imposition can decrease a foundation's impact dramatically, from a patient perspective, and patient actors in many circumstances should be willing to pay large fractions of their budgets to avoid it.

There are three primary ways for a philanthropist to retain the freedom to spend on a patient schedule, each with its strengths and weaknesses.

First, she may establish her foundation in a country that does not impose a disbursement minimum, such as the United Kingdom.

As alluded to in the previous subsection, however, though tax regimes (including in the UK's) commonly allow for tax-exempt trusts established to pursue precise charitable objectives, the US is unusually liberal in allowing tax-exempt foundations that engage in philanthropy at the wide discretion of their managers. Furthermore, and perhaps relatedly, philanthropy is a much bigger presence in the US than in any other developed country.[66] Recall that the US Congress did not bother to impose a disbursement requirement until 1969: 56 years after the pioneering establishment of the Rockefeller Foundation, and as perpetual foundations were just beginning to become a mainstay of the American nonprofit landscape. Canada—another unusually philanthropic country—followed suit in 1976.[67] Other countries may thus impose a similar requirement if big, let alone patient, philanthropy ever takes root there.

Second, she may give to a donor-advised fund (DAF). DAFs are foundations to which others may donate, and pay (usually small) fees, in exchange for the right to "advise" on how these donations are invested and on when and to what nonprofits they are ultimately disbursed. DAFs in turn have a reputational incentive to ensure that the donor's advice is always followed. DAF

[66] This can be quantified in many ways, but see e.g. Charities Aid Foundation (2016) for international data on charitable giving by individuals as a percentage of GDP.

[67] https://www.canada.ca/en/department-finance/programs/consultations/2021/boosting-charitable-spending-communities/backgrounder-disbursement-quota-consultation.html

foundations as a whole are subject to standard disbursement requirements. However, because the contributors to large existing DAFs collectively disburse 20–25% of their assets per year on average (Andreoni, 2018), any individual contributor does not face a disbursement minimum, and in principle generally will not until her assets grow to at least 75% of the assets of the DAF foundation in question (at which point 20% of its remaining assets would constitute less than 5% of the total). With the world's largest DAF—Fidelity Charitable—holding $35b under management as of 2020 (Fidelity Charitable, 2020), and American DAFs as whole holding $142b in 2019 (National Philanthropic Trust, 2020), DAFs might appear able to absorb a considerable quantity of patient, tax-exempt funding without the difficulties of establishing a new, and acceptably flexible, foundation overseas.

Givers to DAFs, however, have no legal right to direct how these funds are invested or regranted. DAFs to date have rarely allocated donor funds not as advised, but if a single DAF account came to constitute a large enough share of some DAF's assets under management, the incentive to expropriate it could outweigh the reputational incentive to manage it faithfully. DAFs also give their account holders relatively little flexibility in investment choices, making it difficult to implement the investment considerations of Sections 4.1–4.3. Finally, as DAFs have come to play a larger role in American philanthropy, calls to impose more stringent disbursement requirements on the sub-industry have intensified: recall for example Schleifer (2019). These calls have culminated in the proposed ACE Act, discussed briefly in Section 1.1, which would require DAF accounts to empty in a limited timeframe. These calls and policy proposals have been directed not only at the largest or fastest-growing DAFs but at DAFs as a whole.

Finally, a patient philanthropist may invest her assets in a non-tax-exempt foundation or US trust, neither of which is subject to disbursement requirements. For a patient philanthropist, the value of flexibility in disbursement scheduling probably, by my estimate, far exceeds the value of the tax benefit to investing within a tax-exempt foundation (though these face only a 1.39% tax on investment income, rather than the top marginal rate of 23.8% levied otherwise). Establishing a tax-exempt American foundation thus appears to be ill-advised for patient philanthropists, except as a pass-through vehicle for holding funds in the years just before disbursement.

However, as in the cases of tax-exempt American and Canadian foundations and American DAFs, the rise of patient philanthropic behavior within taxed institutions may ultimately motivate more stringent regulation or taxation of these institutions as well. Just as increases in income or wealth inequality are often met with calls for redistribution—even if the fortunes driving the inequality were also taxed as they grew—the growth of large and patient philanthropic institutions, however taxed, would likely generate calls for further taxation or mandatory disbursement. Given this risk, whether it is worthwhile to make a long-term philanthropic investment plan depends, as in the discussion of value drift risk in Section 5.2, on whether the benefits of delaying spending grow monotonically until the patient-optimal spending time. If they do, the risk of future legal obstacles to patient behavior does not affect the sign of the value of delaying spending from the present. Otherwise, these obstacles may render it

preferable to spend today than to delay: they may imply that the value of delaying seemed positive only because of a high value placed on the option to spend on what is actually an infeasible timescale.

In short, the magnitude, and perhaps even the sign, of the value of patient philanthropic behavior depends on the long-term ability of philanthropic institutions *as a class* to grow for centuries without their collective size or parsimony triggering a punitive tax or regulatory response. The history of the waqf system suggests that this is possible; the more recent history of philanthropy in the Western world suggests that it will be difficult. One ground for optimism, however, is that patient philanthropists genuinely value the future more than the populations and governments with whom they must continually contend. In principle, therefore, there should be a price that the patient are willing to pay for the right to grow unmolested in centuries to come, which impatient policymakers today are willing to accept. This transaction does not take place only because an impatient policymaker does not have a power he wishes he had, namely the power to commit himself and his successors to laissez-faire disbursement policy in exchange for an up-front transfer. There are in principle dynamic mechanisms that can render it never profitable for the impatient to renege, and exploring these is a subject of ongoing research, including my own. But for now let us simply emphasize that, if one can be found and implemented, all parties will consider it to their benefit.

5.4 *The success of the patient*

Another ground for optimism is a proliferation of recent evidence that patient parties do indeed typically save more, and thereby successfully grow more quickly, than their less patient counterparts. The argument that those who care more about the future should and can do better by spending more slowly, therefore, has at least some direct empirical support.

Epper et al.'s (2020) longitudinal study of Danish adults finds that those who exhibit more patience in an experimental setting grow richer, over the subsequent 15 years, than their less patient peers. The ⅓ most patient subjects are consistently positioned six percentage points higher in the wealth distribution than the ⅓ least patient, and this can largely be attributed to the patient exhibiting higher saving rates.

Of course, these gaps can grow more substantial when transmitted intergenerationally. Individual time preferences are far from perfectly heritable, but substantial, time-preference-based differences in growth rates across households in an intergenerational setting are still possible when average time preferences differ geographically. To explore this relationship at a larger scale, therefore, Sunde et al. (2021) study the cross-sectional relationship between patience on the one hand, and income and capital accumulation on the other, across nations, sub-national regions, and individuals. (Their data on patience for the sample comes from a survey, but its responses have been experimentally validated as genuinely tracking time preferences.) They find that patience explains over 30% of cross-country variation in income, even after controlling for a host of other factors, and that the impact of patience can be accounted for by the extent to which the patient accumulate more capital and more years of education. They find similar, though smaller, relationships between patience and income at the individual and sub-national region levels as well.

Finally, and most ambitiously, Schmelzing (2020) documents that (European and then global) risk-free interest rates have been declining since at least 1300. Stefanski and Trew (2020) argue that the only explanation for this decline compatible with the data is a decline in investor time preference. They then calculate that this decline can in principle be entirely explained by the increasing concentration of capital in patient hands. Interestingly, Stefanski and Trew do not assume that any wealth is transmitted intergenerationally. Instead, they use a quasi-Malthusian model in which wealth patiently accumulated within a generation is spent on fertility, and they calibrate the model to existing research on the observed distribution of time preference (as reflected in savings behavior), the heritability of time preference, and the historical relationship between wealth and fertility. Still, the model requires the patient in each generation to accumulate more capital than the impatient, and allowing for bequests would presumably generate similar results.[68]

As with the analyses of institutional longevity in Section 5.1, any attempt to draw philanthropic lessons from these individual-, country-, and economy-level findings can only be speculative. The history of the waqf system tells us that a nonprofit sector as a whole can grow to vast proportions. Until more research is done on the historical determinants of this growth and on how it varied by institution, however, we cannot turn to waqfs for conclusions or even case studies about the long-term performance of patient philanthropic financial behavior. Any such conclusions would also, of course, be drawn from a particular time and culture, and may not apply to philanthropy in the Western world today.

Thankfully, there are relatively successful examples of more recent, and perhaps more culturally relevant, implementations of patient philanthropy. Upon his death in 1790, for instance, Benjamin Franklin—almost, it seems, on a whim—established trusts for the cities of Boston and Philadelphia. Each was endowed with a sum equivalent to no more than $200k in today's dollars. They were to be disbursed according to the cities' wishes in two installments, in 1890 and 1990 respectively. In the meantime, they were to be lent at 5% interest to young local tradesmen looking to start their careers. The trusts did not grow as quickly as intended, in part because as the economy changed, it eventually grew difficult to find local tradesmen who met Franklin's criteria and who were interested in loans made out to Franklin's specifications. Nevertheless, the trusts ultimately produced gifts of about $6.5m and $1.5m to the two cities after the first century, in present-value terms, and of $5m and $2.25m after the second.[69]

Suppose consumption per capita in the two cities from 1790 to 1890 and from 1890 to 1990 grew by the same proportion as GDP per capita in the United States as a whole. Suppose also that the welfare benefits of the schools and museums the trusts ultimately funded (or the projects Franklin would otherwise have funded in 1790) were a constant multiple, in each year, of the welfare benefits of consumption in that year. Finally, suppose that the loans to the tradesmen had no benefit beyond the interest that accrued to the trusts. Under these assumptions, I estimate that the welfare impact of Franklin's investment plan was about 3x what

[68] By comparing the annualized interest rates implied by the prices of long-term leases on real estate, it appears that interest rates are forecast to continue falling in the future as well (Giglio et al., 2015). A forecast of declining discount rates, however, could be generated by many hypotheses beyond the hypothesis that the average investor will grow more patient.

[69] https://chriscarosa.com/2017/08/ben-franklin-trusts-did-they-work/; https://measuringworth.com/calculators/uscompare/; https://measuringworth.com/datasets/usgdp/

the welfare impact of spending in 1790 would have been. Put another way, I estimate that as long as the city governments of Boston and Philadelphia, in 1890 and 1990, were at least ⅓ as good as Franklin was in 1790 at spending funds on valuable public projects, Franklin made the right decision from a patient perspective. This benefit ratio, while significant, is smaller than the benefit ratios promised in Section 3.3 for several transparent reasons, including the restrictive lending policies discussed above and the failure to invest on an even longer time horizon. After learning from his example, and with the benefit of the more extensive thought that has been put into this question since his day, it seems likely that a patient philanthropist today could do even better, and by a substantial margin.

Philanthropic trusts charged with investing for long periods before disbursing have been established on at least a few other relatively recent occasions. Of those that survived immediate legal challenges, the least successful of which I am aware was created by Anna C. Mott. Mott endowed a trust in 1902 with a sum equivalent to about $25k today, to be invested for one hundred years and then given to the city of Toledo, Ohio. The trust earned a paltry 2.5% real rate of return, and in 2002, it paid out a sum equivalent to about $300k today.[70] Under assumptions analogous to those above, I estimate that the welfare impact of this plan was only negligibly (1.06x) higher than the welfare impact of spending in 1902 would have been.

That said, I know of no such trusts that appear to have performed *worse*, on these rough calculations, than immediate giving.[71] The abject failures all appear to have been family trusts intended for distant descendants, or long-term endowments for philanthropic purposes whose stewardship was nonetheless entrusted to family members, which the family members promptly sued and won the right to seize for themselves.[72] (Indeed, the former are now outlawed in much of the world, including the UK and most of the US.)

Though the performance of university endowments has been studied extensively, as noted in Section 4.3, I do not believe there has not yet been a systematic analysis of the performance of philanthropic institutions so patient as to invest a century before disbursing. For our purposes, of course, this analysis would be at least somewhat valuable. It would probably not produce definitive conclusions, however, because so few of these institutions have been established, and none yet seem to have been very large or very thoughtfully executed.[73] Given this unavoidable uncertainty, though, the case for attempting long-term philanthropic investment is strengthened by an asymmetry in the stakes involved. In light of the analysis presented throughout this report, I would argue that it is at least plausible that a successful implementation could do many

---

[70] Scott (2002); https://measuringworth.com/calculators/uscompare/; https://measuringworth.com/datasets/usgdp/

[71] These calculations assume that the cities (and other parties whose behavior affected the economic trajectories of the cities) did not substantially alter their own spending behavior through the centuries in light of the coming transfers. That is, they assume that the philanthropists' spending schedules can be evaluated according to what I have called "BP". If the cities responded fully "strategically", Mott's trust underperformed an immediate transfer, and Franklin's case is ambiguous.

[72] See Collins (2011) for two examples.

[73] A new patient philanthropic trust—moderately larger than Franklin's and Mott's, (I hope) more thoughtfully executed, and far more flexible—has recently been established in London. Named the Patient Philanthropy Fund, it was motivated in large part by the reasoning presented throughout this document. I am a member of its management committee. If nothing else, it may offer one more data point for researchers in future centuries.

times more good than philanthropy in the present day. Even an entirely failed implementation, on the other hand, is unlikely to do much worse than to have no impact at all.

5.5 *A patient world*

As noted at the beginning of the previous subsection, the world appears to be getting more patient over time. That is, there is evidence that the time preference reflected in interest rates has gradually been falling. To be sure, the evidence for this purported millennium-long historical trend is far from watertight. As long as households and institutions with different time preferences remain free to spend and save at their own preferred rates, however, it is hard to escape the conclusion that a trend along these lines will unfold in the future.

The discussion of this document so far has assumed that the value of spending patiently, for a patient philanthropist, consists entirely in the value of directly shifting the schedule of spending on his chosen causes. When we reflect on the above dynamic, however, we see that even more is at stake. Capital accumulation lowers the gap between interest rates and growth rates, and by the mechanisms described in Sections 2.2 and 2.3, this in turn raises the effective patience of society as a whole, including its firms and governments.

A patient philanthropist's impact on growth and interest rates would be slight even if these variables could be moved without altering others' behavior. In practice, moreover, we should expect these movements to be partially offset by impatient households saving less as the interest rate falls. At the same time, the corporate and public budgets allocated slightly more patiently because of this impact are vast. The importance of impacts through this channel, relative to the importance of the more direct impacts of a philanthropist's choice of disbursement schedule, depends on many variables we have not considered here, such as the extent to which the philanthropist's spending preferences differ from those of policymakers and typical citizens along dimensions other than time preference. For reference, however, if they do not differ at all, I roughly estimate that from the philanthropist's perspective, the importance of the two classes of impacts are of the same order of magnitude.

In short, this happy side-effect of patient financial behavior is not inconsequential. By squeezing a drop of impatience out of the interest rate all parties face, patient investors substantively reshape the rest of society along faster-growing, more future-oriented lines.

This is more than theoretical speculation. The effect of patient investment on interest rates is widely held to be among the primary causes of the most important event in economic history.

As Allen (2009) argues at length, many of the labor-saving devices responsible for the first stages of the Industrial Revolution, such as the spinning jenny, were obvious long before they were adopted. They were adopted in England in the 1700s only because this was the first time and place in which the cost of capital (i.e. the interest rate) had grown low enough, and the cost of labor (the wage rate) high enough, to render their adoption profitable. In other words, they were adopted because of the increase to the capital-to-labor ratio in northwestern Europe that immediately predated the Industrial Revolution and did not occur elsewhere. That increase, in

turn, can straightforwardly be attributed at least in part to the fact (Humphries and Weisdorf, 2019) that, throughout the preceding centuries, people in northwestern Europe worked harder and saved more.

This last phenomenon has been noted, and attributed to a variety of causes, at least since Max Weber's hypothesis of a "Protestant work ethic". For our purposes, however, the ultimate attribution is irrelevant. What matters is that an episode of unusually intense capital accumulation, whatever its source, appears likely to have motivated profit-seeking manufacturers to invest in more technological advances than they would have otherwise. On this account, the English workers and savers of the 17th century did not simply grow their family fortunes at a faster clip. They also unconsciously incentivized other actors in the economy around them to adapt, and contribute, to a regime of faster growth.

The most important benefit of lowering the impatience reflected in interest rates in, say, 21st century America may be very different from the most important in 18th century England. Lowering the time preference reflected in interest rates would lead governments to devote more resources to monitoring asteroids or combating climate change, for instance; as argued in Section 3.4, policy changes along these lines are probably now more important from a patient perspective than speeding growth by incentivizing the development of labor-saving technology. But the force at work would be similar.

On the other hand, suppose that others on the whole will grow more patient regardless of what we do. Then the rate at which we can do more good by delaying our spending—either by better allocating resources on our chosen causes across time, or by the second channel discussed above—is slowly diminishing. We may be able to do, say, 3% more good by investing a dollar to spend on a given cause next year than by spending it today, but we should expect that the gain from investing it from next year to the following year will be slightly less than 3%. If the estimates from Stefanski and Trew are to be believed, and if the trend they identify can be projected into the future, there appear to be only a few centuries left before the time preference reflected in interest rates is negligible.

The rise of patience will be an unambiguously positive development, from our perspective, if the newly patient actors all have preferences aligned with our own (except for their former impatience). In this case, however, the prevalence of others' patient behavior will shrink the relative value of our own patient behavior. It may even mean that we do better to begin spending today when we would otherwise have done best to delay for some period. This is simply a larger-scale case of the observation from Section 3.5 that a community of patient philanthropists is effectively "richer", and so should spend more quickly, if there will be new sources of patient funding on its favored causes in the future.

Of course, however, many actors may have preferences more deeply misaligned with one's own. We have not focused on these "conflictual" dynamics in this document, but they may be highly important, especially over the long term. The prospect of less friendly actors growing more patient is not a hope but a threat, and it should motivate not faster spending but more urgency around plans for long-term investment.

At present, because of the time preference reflected in interest rates, the future is in effect cheaply for sale. Consider, for illustration, the price difference between a land freehold and a 100-year lease: the latter typically costs almost 90% as much as the former (Giglio et al., 2015). Buying land, and then selling a long-term lease on it, thus currently gives the buyer the right to determine what is done with the land from year 2121 until our civilization collapses or the world ends, almost for free. Once a sufficiently large proportion of the world's resources are being invested along patient lines, opportunities of this kind will no longer be available. The future will have been bought. If its owners do not put it to good use, this will be a massive and irrevocable loss.

In short, a patient world is coming. When it comes, as long as nothing goes too disastrously, it will be here to stay. As philanthropists, the plans we make today stand a good chance of influencing how quickly it comes and what shape it takes. In Section 3.6, we considered the possibility that the present offers patient philanthropists fleeting and uniquely impactful opportunities for spending. Even if this is true, it must be weighed against the possibility that the present offers patient philanthropists a fleeting and uniquely impactful opportunity to invest. Once again, therefore, we see that the philanthropist's disbursement scheduling problem, despite its complexity, is worth engaging with as seriously as possible. By playing the game well, patient philanthropists may be able not only to allocate their own spending more efficiently over time—which could be very valuable in its own right—but to have a substantial impact on the course of the future.

Bibliography

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2001. "The Colonial Origins of Comparative Development: An Empirical Investigation." *American Economic Review* 91, no. 5: 1369–1401. [Link]

Adiguzel, Fatih Serkant, and Timur Kuran. 2021. "The Islamic Waqf: Instrument of Unequal Security, Worldly and Otherworldly", ERID Working Paper 305. Economic Research Initiatives at Duke. [Link]

Alan, Sule, and Martin Browning. 2010. "Estimating Intertemporal Allocation Parameters Using Synthetic Residual Estimation." *The Review of Economic Studies* 77, no. 4: 1231–1261. [Link]

Alexander, Scott. 2013. "Investment and Inefficient Charity." [Link]

Allen, Robert C. 2009. *The British Industrial Revolution in Global Perspective*. Cambridge University Press.

Andreoni, James. 2018. "The Benefits and Costs of Donor-Advised Funds." *Tax Policy and the Economy* 32, no. 1: 1–44. [Link]

Aschenbrenner, Leopold. 2020. “Existential Risk and Growth”, *GPI Working Paper* 6-2020. Global Priorities Institute. [Link]

Azlen, Michael W., and Ilan Zermati. 2017. “Investing Like the Harvard and Yale Endowment Funds.” [Link]

Barrage, Lint. 2018. “Be Careful What You Calibrate for: Social Discounting in General Equilibrium.” *Journal of Public Economics* 160: 33–49. [Link]

Bartsch, Sarah M., Peter J. Hotez, Lindsey Asti, Kristina M. Zapf, Maria Elena Bottazzi, David J. Diemert, and Bruce Y. Lee. 2016. “The Global Economic and Health Burden of Human Hookworm Infection.” *PLoS Neglected Tropical Diseases* 10, no. 9: e0004922. [Link]

Bilgin, Doga, and Reinhard Ellwanger. 2019. *The Simple Economics of Global Fuel Consumption*. Bank of Canada. [Link]

Breeden, Douglas T. 1979. “An Intertemporal Asset Pricing Model With Stochastic Consumption and Investment Opportunities.” *Journal of Financial Economics* 7, no. 3: 265–296. [Link]

Bressler, R. Daniel. 2021. “The Mortality Cost of Carbon.” *Nature Communications* 12, no. 1: 1–12. [Link]

Callahan, David. 2018. *The Givers: Wealth, Power, and Philanthropy in a New Gilded Age*. Vintage.

Campbell, John. 2010. “Investing and Spending: The Twin Challenges of University Endowment Management.” In *Forum for the Future of Higher Education*. [Link]

Carosa, Chris. 2017. “Ben Franklin Trusts – Did They Work?” [Link]

Center for Strategic Philanthropy and Civil Society. 2007. “The Green Revolution: Rockefeller Foundation, 1943.” [Link]

Chambers, David, Elroy Dimson, and Charikleia Kaffe. 2020. “Seventy-Five Years of Investing for Future Generations.” *Financial Analysts Journal* 76, no. 4: 5–21. [Link]

Chang, Wayne. 2020. “Comment on ‘The Case for Investing to Give Later’”[Link]

Charities Aid Foundation. 2016. *Gross Domestic Philanthropy: An International Analysis of GDP, Tax and Giving*. [Link]

Charity Navigator. 2018. “Giving Statistics.” [Link]

Chen, Ariel, Kathryn H. Jacobsen, Ashish A. Deshmukh, and Scott B. Cantor. 2015. “The Evolution of the Disability-Adjusted Life Year (DALY).” *Socio-Economic Planning Sciences* 49: 10–15. [Link]

Chetty, Raj. 2006. “A New Method of Estimating Risk Aversion.” *American Economic Review* 96, no. 5: 1821–1834. [Link]

Christiano, Paul F. 2019. “Analyzing Divestment.” [Link]

Çizakça, Murat. 1993. Risk Sermayesi Özel Finans Kurumlarıve Para Vakıfları. İslami İlimler Araştırma Vakfı (İSAV). [Link]

— . 2004. “Ottoman Cash Waqfs Revisited: The Case of Bursa 1555–1823.” [Link]

Collins, Paul. 2011. “Trust Issues.” *Lapham’s Quarterly*. [Link]

Copenhagen Consensus Center. 2015. “Post-2015 Consensus.” [Link]

Cotton-Barratt, Owen. 2015. *Allocating Risk Mitigation Across Time*. Tech. Rep. 2015-2. Future of Humanity Institute [Link]

Council of Economic Advisers to President Obama. 2017. “Discounting for Public Policy: Theory and Recent Evidence on the Merits of Updating the Discount Rate.” [Link]

Council on Foundations. 2021. “Council on Foundations Opposes the Accelerating Charitable Efforts Act.” [Link]

Cowen, Tyler. 1992. “Consequentialism Implies a Zero Rate of Intergenerational Discount.” In *Philosophy, Politics, and Society: Volume 6, Justice Between Age Groups and Generations*. Ed. by Peter Laslett and James S. Fishkin. Yale University Press. [Link]

— . 2006. *Good and Plenty: The Creative Successes of American Arts Funding*. Princeton University Press.

Cowen, Tyler, and Derek Parfit. 1992. “Against the Social Discount Rate.” In *Philosophy, Politics, and Society: Volume 6, Justice Between Age Groups and Generations*. Ed. by Peter Laslett and James S. Fishkin. Yale University Press. [Link]

Daepp, Madeleine I. G., Marcus J. Hamilton, Geoffrey B. West, and Luís M. A. Bettencourt. 2015. “The Mortality of Companies.” *Journal of The Royal Society Interface* 12, no. 106: 20150120. [Link]

Della Porta, Donatella, and Mario Diani, eds. 2016. *The Oxford Handbook of Social Movements*. Oxford University Press.

Dickens, Michael. 2019. “Should Global Poverty Donors Give Now or Later?” [Link]

— . 2020. “How Much Leverage Should Altruists Use?” [Link]

— . 2021. “Uncorrelated Investments for Altruists.” [Link]

Dillon, Charles. 2021. “An Analysis of Metaculus Predictions of Future EA Resources, 2025 and 2030.” [Link]

Dimon, Jamie, and Warren E. Buffett. 2018. “Short-Termism Is Harming the Economy.” *The Wall Street Journal* (6 June 2018). [Link]

Drupp, Moritz A., Mark C. Freeman, Ben Groom, and Frikk Nesje. 2018. “Discounting Disentangled.” *American Economic Journal: Economic Policy* 10, no. 4: 109–34. [Link]

Epper, Thomas, Ernst Fehr, Helga Fehr-Duda, Claus Thustrup Kreiner, David Dreyer Lassen, Søren Leth-Petersen, and Gregers Nytoft Rasmussen. 2020. “Time Discounting and Wealth Inequality.” *American Economic Review* 110, no. 4: 1177–1205. [Link]

Federal Highway Administration. *Highway Safety Benefit-Cost Analysis Guide*. 2018. U.S. Department of Transportation. [Link]

Fidelity Charitable. 2020. *2020 Annual Report*. [Link]

Fleishman, Joel L., J. Scott Kohler, and Steven Schindler. 2009. *Casebook for the Foundation: A Great American Secret*. PublicAffairs.

Flood, Chris. 2019. “Leading US University Endowments Struggle to Beat Tracker Funds.” *Financial Times* (8 Dec. 2019). [Link]

Frederick, Shane, George Loewenstein and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: a Critical Review." *Journal of Economic Literature* 40, no. 2: 351–401. [Link]

Freeman, Larry. 2020. "How Long Did the Empires of Ancient Civilizations Last?" [Link]

Freitas-Groff, Zachary. 2020. "Climate Warriors and Deficit Hawks: Eliciting Intergenerational Discount Rates."

Gandelman, Néstor, and Rubén Hernández-Murillo. 2015. "Risk Aversion at the Country Level." *Federal Reserve Bank of St. Louis Review* 97, no. 1: 53–66. [Link]

Gates, Bill. 2017a. "Mosquito Wars." [Link]

— . 2017b. "Our Education Efforts Are Evolving." [Link]

Geczy, Christopher, Jessica S. Jeffers, David K. Musto, and Anne M. Tucker. 2021. "Contracts With (Social) Benefits: The Implementation of Impact Investing." *Journal of Financial Economics* . [Link]

Gieczewski, Germán. 2021. "Policy Persistence and Drift in Organizations." *Econometrica* 89, no. 1: 251–279. [Link]

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2015. "Very Long-Run Discount Rates." *The Quarterly Journal of Economics* 130, no. 1: 1–53. [Link]

Gilbert, Thomas, and Christopher Hrdlicka. 2015. "Why Are University Endowments Large and Risky?" *The Review of Financial Studies* 28, no. 9: 2643–2686. [Link]

Giving USA. 2021. *Giving USA 2021: The Annual Report on Philanthropy for the Year 2020*. [Link]

Government of Canada. 2020. *Backgrounder for Disbursement Quota Consultation*. [Link]

Gunn-Graffy, Colette. 2020. "History of the Scholarship." The Rhodes Project. [Link]

Hager, Sandy Brian, and Joseph Baines. 2020. "Biggest Companies Pay the Least Tax, Leaving Society More Vulnerable to Pandemic." [Link]

Hillebrandt, Hauke, and John Halstead. 2018. *Impact Investing Report*. Founders Pledge. [Link]

HM Treasury. 2020. *The Green Book*. [Link]

Hoeijmakers, Sjir. 2020. "Investing to Give." Founders Pledge. [Link]

Holcombe, Randall G. 2000. *Writing Off Ideas: Taxation, Foundations, and Philanthropy in America*. Routledge.

House of Lords Select Committee on Charities. 2016. *Response to Call for Written Evidence: Further Evidence*. The Association of Charitable Foundations. [Link]

Humphries, Jane, and Jacob Weisdorf. 2019. "Unreal Wages? Real Income and Economic Growth in England, 1260–1850." *The Economic Journal* 129, no. 623: 2867–2887. [Link]

Hurford, Peter. 2013. "Giving Now Currently Seems to Beat Giving Later." [Link]

Inauen, Emil, Katja Rost, Margit Osterloh, and Bruno S. Frey. 2010. "Back to the Future: A Monastic Perspective on Corporate Governance." *Management Revue*: 38–59. [Link]

Initiative to Accelerate Charitable Giving. 2021. "Common-Sense Reforms." [Link]

Johnston, Louis, and Samuel H. Williamson. 2021. “What Was the U.S. GDP Then?” [Link]

Kalai, Adam, and Ehud Kalai. 2001. “Strategic Polarization.” *Journal of Mathematical Psychology* 45, no. 4: 656–663. [Link]

Kelly, J. L., Jr. 1956. “A New Interpretation of Information Rate.” *Bell System Technical Journal* 35, no. 4: 917–926. [Link]

Koller, Tim, James Manyika, and Sree Ramaswamy. 2017. “The Case Against Corporate Short Termism.” [Link]

Kuran, Timur. 2001. “The Provision of Public Goods Under Islamic Law: Origins, Impact, and Limitations of the Waqf System.” *Law and Society Review*: 841–898. [Link]

Leat, Diana. 2016. *Philanthropic Foundations, Public Good and Public Policy*. Palgrave Macmillan.

Levine, Joseph, Maximilian Negele, and Philip Trammell. In Progress. “The Longevity of Nonprofit Institutions.”

Lockwood, Lee M. 2012. “Bequest Motives and the Annuity Puzzle.” *Review of Economic Dynamics* 15, no. 2: 226–243. [Link]

MacAskill, William. 2019. “Are We Living at the Most Influential Time in History?” [Link]

— . 2020. “Are We Living at the Hinge of History?”, *GPI Working Paper* 12-2020. Global Priorities Institute. [Link]

McAdam, Doug, and David A. Snow. 2010. *Readings on Social Movements: Origins, Dynamics and Outcomes*. Oxford University Press.

Mercier, Guillaume, and Ghislain Deslandes. 2017. “There Are No Codes, Only Interpretations: Practical Wisdom and Hermeneutics in Monastic Organizations.” *Journal of Business Ethics* 145, no. 4: 781–794. [Link]

Merton, Robert C. 1993. “Optimal Investment Strategies for University Endowment Funds.” In *Studies of Supply and Demand in Higher Education*, ed. by Charles T. Clotfelter and Michael Rothschild. University of Chicago Press. [Link]

Metaculus. 2021. “How Much Will GiveWell Guess It Will Cost to Get an Outcome as Good as Saving a Life, at the End of 2031?” [Link]

Morris, Peter. 2018. “The Aim Is Not the Same.” [Link]

NASA. 2021. “Planetary Defense Coordination Office.” [Link]

National Philanthropic Trust. 2020. *The 2020 DAF Report*. [Link]

Negele, Maximilian. 2020. “Delegation of Irreversible Decisions Under Uncertainty.” MPhil thesis, University of Oxford.

Newberry, Toby. 2021. “How Cost-effective Are Efforts to Detect Near-Earth Objects?” *GPI Technical Report* T1-2021. Global Priorities Institute. [Link]

Nordhaus, William D. 2009. “An Analysis of the Dismal Theorem.” *Cowles Foundation Discussion Paper* 1686. Cowles Foundation for Research in Economics. [Link]

— . 2004. “Schumpeterian Profits in the American Economy: Theory and Measurement”, *NBER Working Paper Series* 10433. NBER. [Link]

OECD. 2020. *Main Science and Technology Indicators*. [Link]

Ord, Toby. 2020. *The Precipice: Existential Risk and the Future of Humanity*. Hachette Books.

— . In Progress. "A Simple Framework for Comparing Longterm Effects."

Oster, Sharon M. 2001. "The Effect of University Endowment Growth on Giving: Is There Evidence of Crowding Out?" [Link]

Parfit, Derek. 1984. *Reasons and Persons*. Oxford University Press.

— . 2011. *On What Matters*. Vol. 2. Oxford University Press.

Pharoah, Cathy, and Jenny Harrow. 2010. "Payout With an English Accent: Exploring the Case for a Foundation 'Distribution Quota' in the UK." *Presented at ARNOVA 39th Annual Conference*. [Link]

QuantStart. 2014. "Money Management via the Kelly Criterion." [Link]

Reich, Rob. 2018. *Just Giving: Why Philanthropy Is Failing Democracy and How It Can Do Better*. Princeton University Press.

Risi, Joseph, Amit Sharma, Rohan Shah, Matthew Connelly, and Duncan J. Watts. 2019. "Predicting History." *Nature Human Behaviour* 3, no. 9: 906–912. [Link]

Robinson, Zack. 2016. "The Donor's Dilemma: The Time Value of Money and Maximizing Effectiveness" [Link]

Rose, Alexander. 2020. "The Data of Long-Lived Institutions." [Link]

Rost, Katja, Emil Inauen, Margit Osterloh, and Bruno S. Frey. 2010. "The Corporate Governance of Benedictine Abbeys: What Can Stock Corporations Learn From Monasteries?" *Journal of Management History*. [Link]

Roth, Benjamin. 2021. "Impact Investing: A Theory of Financing Social Enterprises." *Harvard Business School Entrepreneurial Management Working Paper* 20–078. [Link]

Roth Tran, Brigitte. 2019. "Divest, Disregard, or Double Down? Philanthropic Endowment Investments in Objectionable Firms." *American Economic Review: Insights* 1, no. 2: 241–56. [Link]

Sagan, Carl. 1994. *Pale Blue Dot: A Vision of the Human Future in Space*. Random House.

Sagan, Carl, and Steven J. Ostro. 1994. "Long-Range Consequences of Interplanetary Collisions." *Issues in Science and Technology* 10, no. 4: 67–72. [Link]

Sandberg, Anders. In Progress. "Grand Futures."

Savoie, Joseph. 2018. "Empirical Data on Value Drift" [Link]

Schleifer, Theodore. 2019. "How a Lawsuit Could Reveal Secrets About Silicon Valley's Favorite Philanthropic Loophole." *Vox*. [Link]

Schmelzing, Paul. 2018. "Eight Centuries of Global Real Interest Rates, R-G, and the 'Suprasecular' Decline", *1311–2018*, *Staff Working Paper* 845. Bank of England. [Link]

Schönholzer, David, and Eric Weese. 2018. "Creative Destruction in the European State System: 1000–1850." [Link]

Scott, Rebekah. 2002. "100-Year-Old $1,000 Gift to Toledo Soars to $213,000 for Parks." *The Blade* (29 Jan.2002). [Link]

Shulman, Carl. 2012. “Salary or Startup? How Do-Gooders Can Gain More From Risky Careers.” [Link]

Smith, Jeremy N. 2017. *Epic Measures: One Doctor. Seven Billion Patients*. HarperOne.

Stefanski, Radoslaw, and Alex Trew. 2020. “Patience Breeds Interest: The Rise of Societal Patience and the Fall of the Risk-Free Interest Rate.” *School of Economics and Finance Discussion Papers* 2002. University of St. Andrews. [Link]

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim. 2021. “Patience and Comparative Development.” [Link]

Taagepera, Rein. 1978a. “Size and Duration of Empires: Systematics of Size.” *Social Science Research* 7, no. 2: 108–127. [Link]

— . 1978b. “Size and Duration of Empires Growth-Decline Curves, 3000 to 600 B.C.” *Social Science Research* 7, no. 2: 180–196. [Link]

— . 1979. “Size and Duration of Empires: Growth-Decline Curves, 600 B.C. to 600 A.D.” *Social Science History* 3, no. 3/4: 115–138. [Link]

The Rockefeller Foundation. 2016. “The Rockefeller Foundation Announces $6 Million in Additional Funding for 100,000 Inner City Students Nationwide to See HAMILTON.” [Link]

— . 2017. “Eradicating Hookworm.” [Link]

The Scottish Hospital of the Foundation of King Charles II. 2019. *Report and Financial Statements for the Year Ended*. The Royal Scottish Corporation. [Link]

Todd, Benjamin. 2017. “The Case for Reducing Existential Risks.” [Link]

— . 2020. “More Empirical Data on ‘Value Drift’”[Link]

— . 2021. “How Are Resources in Effective Altruism Allocated Across Issues?” [Link]

Tomasik, Brian. 2021. “Should Altruists Leverage Investments?” [Link]

Trammell, Philip. 2021. “Dynamic Public Good Provision Under Time Preference Heterogeneity: Theory and Applications to Philanthropy”, *GPI Working Paper* 9-2021. Global Priorities Institute. [Link]

— , and Nuño Sempere. In Progress. “Labor, Capital, and Patience in the Optimal Growth of Social Movements.” [Link]

U.S. Bureau of Labor Statistics. 1915. *Report of the Commission on Industrial Relations*. [Link]

U.S. Office of Management and Budget. 2019. *Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*. [Link]

UNESCO. 2020. “How Much Does Your Country Invest in R&D?”[Link]

Waleson, Heidi. 2017. Atlantic Insights: Giving While Living. The Atlantic Philanthropies. [Link]

Williamson, Samuel H. 2021. “Seven Ways to Compute the Relative Value of a U.S. Dollar Amount, 1790 to Present.” [Link]