

The Option Value of Life

Susanne Burri*

December 19, 2019

Abstract

This paper argues that under conditions of uncertainty, there is frequently a positive option value to staying alive when compared to the alternative of dying right away. This value can make it prudentially rational for you to stay alive even if it appears highly unlikely that you have a bright future ahead of you. Drawing on the real options approach to investment analysis (see e.g. Dixit and Pindyck, 1994), the paper explores the conditions under which there is a positive option value to staying alive, and it draws out important implications for the problems of suicide and euthanasia.

Keywords: Real options approach, uncertainty, Deprivation Account of the Badness of Death, suicide, euthanasia.

*London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK.
Email: s.burri@lse.ac.uk. URL: <http://www.lse.ac.uk/cpnss/people/susanne-burri>.

*I dwell in Possibility—
A fairer House than Prose—
More numerous of Windows—
Superior—for Doors—*

Emily Dickinson

1 Introduction

Suppose that you visit your physician for an annual check-up. Instead of being told that you are healthy, you are diagnosed with a condition that is both progressive and incurable. Your physician tells you that it is generally quite difficult to manage the symptoms associated with your condition, and that you will most likely suffer persistent and debilitating pain over the years to come. But the physician also tells you that for some patients who are diagnosed with the condition, their symptoms are only ever very mild, even as the disease progresses.

To receive such a diagnosis would certainly be devastating. But once you had gotten over the initial shock, what value would it be appropriate for you to assign to staying alive? In light of the diagnosis received, might it be a good thing for you if you died today? It seems that for most of us, staying alive remains intuitively preferable to dying right away. But it is not clear that such a preference would be prudentially rational. After all, it seems rather unlikely that you have a bright future ahead of you.

My aim in this paper is to explain why a preference for staying alive can be prudentially rational even in cases where your outlook is quite bleak. More specifically, I argue that in the example just given, it would be wrong for you to assess your options by simply comparing the expected value of staying alive with the expected value of dying right away, while assuming that staying alive will most likely bring a considerable amount of pain. Drawing on the real options approach to investment analysis (see e.g. Dixit and Pindyck, 1994), I show how staying alive tends to come with a positive option value that you would miss out on if you died right away. This option value can make staying alive prudentially rational even if, on balance, you do not expect that staying alive will make

your life go better. As long as suicide or euthanasia are readily available, and as long as they allow you to keep in check the potential downsides of staying alive, it is likely that dying today is not in your self-interest.

It is not currently widely recognised that staying alive can come with a positive option value. I argue that it ought to be, as the option value of life provides an important reason in favour of making available not only the option of physician-assisted suicide, but also the options of voluntary and non-voluntary euthanasia. Understanding what circumstances give rise to option value can moreover help us establish whether a person's preference to end their life is rational or irrational, thus enabling us to limit paternalistic interventions aimed at preventing suicide to cases where such interventions are likely to be justified.

My argument proceeds as follows. In section 2, I introduce the Deprivation Account of the Badness of Death, or Deprivation Account, for short. Its key claim is that a person's death is bad for her if she would have lived a better life had she not died this particular death. Under conditions of certainty, we cannot use the Deprivation Account to make sense of the idea that it is preferable to stay alive even if your outlook is rather bleak. Suppose that you know for sure that staying alive will not add to the value of your life. If so, it is not bad for you to die today according to the Deprivation Account, as you will not live a better life if you remain alive. Frances Kamm (1993) argues that this makes the Deprivation Account importantly incomplete. In her view, a complete account of the badness of death needs to be able to explain why it can be preferable for someone to stay alive even if doing so does not make the person's life go better. In section 3, I sketch Kamm's thoughts on the issue, and explain why they are ultimately unconvincing. Kamm suggests that the state of being dead might be intrinsically bad for a person. I show that even if we grant this questionable assumption, it does not necessarily help us make sense of the idea that it can be in a person's self-interest to remain alive when doing so does not make her life go better. In section 4, I argue that once we move to conditions of deliberative uncertainty we can use the Deprivation Account to explain why it can be in someone's self-interest to stay alive even if this does not, on expectation, make her life go better. In a nutshell, once we move to conditions of deliberative uncertainty we can use the

insights of the Deprivation Account only if we combine them with an appropriate decision rule that specifies how to take into account the presence of uncertainty. I contend that the real options approach to investment analysis yields such a rule. Using the real options approach to capture the value of staying alive means modelling death as an *irreversible disinvestment from life*, and staying alive as the choice to keep one's options open by *not disinvesting right away*. Section 4 explores the conditions under which staying alive comes with a positive option value. In section 5, I conclude by considering some of the practical implications of the option value of life.

2 The Deprivation Account of the Badness of Death

The Deprivation Account of the Badness of Death is the currently prevailing explanation of what can make a person's death bad for her.¹ On the assumption that we cease to exist when we die, the Deprivation Account says that a person's death is bad for her if she would have lived a better life² had she not died this particular death. Analogously, a person's death is good for her if she would have lived a worse life had she not died this particular death. The basic thought behind the Deprivation Account is that even if the *state of being dead* cannot be good or bad for a person, a person's death can nevertheless be good or bad for her in virtue of what it prevents from happening. It follows that on the Deprivation Account, we cannot know whether a person's death was good or bad for her unless we know how comparatively valuable a life her death precluded. Under most circumstances, this is difficult to gauge (see McMahan, 2002: 103-35; Bradley, 2009: 47-60). Consider the introductory example. Once you are diagnosed with a progressive and incurable condition, it seems plausible that dying today might be good for you. If continuing to live means suffering terrible pain, then it may well be true that the shorter

¹The basic ideas behind the Deprivation Account go back to Nagel (1970). For refinements of the account, see e.g. Feldman (1992), Broome (2004), or Bradley (2009).

²Both the Deprivation Account and my arguments in this paper are consistent with a wide range of theories about what makes a person's life go well. For reasons of simplicity, I limit my attention to theories that allow us to measure the goodness or badness of different lives on a cardinal scale.

life that you will have lived if you die today is all things considered more valuable than the longer life that you will eventually have lived if you continue to stay alive. You do not, however, have any guarantee that dying today is good for you. After all, it is possible that you would remain largely symptom-free as the disease progresses, in which case dying today would most likely be bad for you.

To fix this epistemic issue, we can construct a case where we simply stipulate what would happen to you if you did not die today. Consider the following:

Two Lives. Imagine that you are given a choice between living two different lives. You can choose either to die painlessly in your sleep tonight, or to remain alive for an additional seven years, in which case you will die painlessly in your sleep seven years from now. If you choose to remain alive, this choice will be erased from your memory, and you will thus not live the remainder of your life knowing the date of your death. There is one small caveat to staying alive: if you choose to go on living, then you know that by the time you die, your life will have the exact same overall value that it would have had had you chosen to die today. On balance, staying alive thus neither adds to, nor does it detract from, the value of your life.

From a self-interested point of view, what choice should you make?³

Two Lives is constructed in such a way that whichever option you choose, you miss out on an alternative life that is no better or worse than the life that you actually live. In this particular case, the Deprivation Account thus recommends an attitude of indifference between the two options that you are presented with.

Two Lives is a curious case. The vast majority of people to whom I have presented the case tell me that they would strongly prefer to live the longer life. But it seems doubtful that such a pronounced preference in favour of living the longer life can be said

³Throughout this paper, my focus is on *self-interested reasons* in favour of staying alive, i.e. on what it is *prudentially rational* for someone to do. Suppose that you choose to stay alive to take care of your ailing mother. This is a self-interested choice only to the extent that you believe that taking care of your mother is good for you. See also footnote 4 below.

to track any self-interested reasons. After all, it is stipulated in *Two Lives* that you will live an equally good life whichever option you choose, and this stipulation seems to leave no room for the possibility that self-interested reasons might speak in favour of one option over the other.⁴

It is possible that those of us who strongly prefer to live the longer life do so based on flawed reasoning, or are guided by forces other than reasoning, such as an ingrained tendency to shy away from death. As for the first alternative, there are many ways in which we might be mistaken about the self-interested reasons that apply to us in *Two Lives*. Suppose, for example, that longer lives are usually better than shorter lives. If this is the case, then it makes sense, heuristically, to prefer to go on living unless you have good reason to think that your future would look very bleak indeed. It is possible that when we are presented with a case such as *Two Lives*, many of us apply a simple such heuristic, even though we are told that in the case at hand, living longer would *not* translate into living a better life. As for the second alternative, it is plausible that those among us who instinctively shy away from death are thereby equipped with an evolutionary advantage, in which case an ingrained tendency to fear death may have become widespread over time.

In my view, a pronounced preference for the longer life in *Two Lives* should indeed be thought to rest on flawed reasoning, potentially combined with an ingrained tendency

⁴This assumes that under certainty, you have a self-interested reason in favour of choosing option *A* over option *B* just in case you will live a better life if you do so. But this is a controversial assumption. Some people argue that from a self-interested perspective, you have reason to take into account how psychologically connected you are to the different temporal stages of yourself to whom different goods and bads accrue. Suppose that in *Two Lives*, if you choose to remain alive, you will first live through 3.5 good years, which will then be followed by 3.5 bad years. On this specification of the case, you may have a self-interested reason in favour of staying alive, as you will most likely be more psychologically connected to your more immediate future self than to your more remote future self. Jeff McMahan (2002, 2019a, 2019b) defends a version of the Deprivation Account, called the Time-Relative Interest Account (TRIA), that takes into account a person's psychological connectedness in this way. While the arguments that I defend in this paper are compatible with McMahan's TRIA, I set aside considerations of psychological connectedness in the main text so as not to unnecessarily complicate my discussion.

to shy away from death. As I see it, when you know for certain that you will not live a better life if you choose to stay alive, it is difficult to defend wanting to stay alive as prudentially rational. In section 4, I will present what I take to be a satisfactory “error theory” that can explain a preference for the longer life. In short, the thought behind this error theory is that a preference to go on living can, and usually will, track self-interested reasons under realistic conditions, even though it does not track such reasons in the highly artificial setup of *Two Lives*.

While I am satisfied with the error theory that I present, Frances Kamm would beg to differ. She thinks that a complete account of the badness of death needs to be able to explain why it could be in our self-interest to prolong our lives even where this is known not to add to the value of our lives. As she puts it, “[s]ince we could prefer to postpone things being all over, even if this did not increase the total amount of goods we had in our life, we must be trying to avoid something about death other than that it diminishes the amount of goods of life we have” (Kamm, 1993: 19). In the next section, I explore how Kamm tries to make sense of what she claims would be a reasonable preference, and I explain why I think that her suggestion fails.

3 Might It Be Bad to Be Dead?

Kamm (1993, 2017) seems to share the intuition that in *Two Lives*, staying alive appears preferable to dying right away. She introduces the figure of the *Limbo Man*, who prefers a longer future over a shorter one even though the longer future holds no more net goods than the shorter one (Kamm, 1993: 49-53). In her view, the fact that we cannot use the Deprivation Account to make sense of the *Limbo Man’s* preference for the longer life indicates that the Deprivation Account is incomplete as an account of what can make a person’s death bad for her. As Kamm (2017: 734) puts it in a more recent contribution, “[o]ne might have concerns about [the Deprivation Account] as an adequate theory of the badness of death” partly because “one might think that, holding constant the beginning point of our life and the assumed good contents of it, from the point of view of concern

with death it would be better for us all to put off the time at which our conscious life will be all over.”

Kamm (1993: 19; 50) suggests the following explanation for the comparative badness of living the shorter life in a case such as *Two Lives*: even assuming that we cease to exist when we die, the *nothingness* that follows our existence could nevertheless be intrinsically bad for us. If this is the case, then it makes sense to prefer to put off this nothingness for as long as possible. Admittedly, whatever intrinsic badness there is to our posthumous non-existence is necessarily non-experiential, and Kamm’s proffered explanation is thus not compatible with theories of value that stipulate that all goodness or badness for a person is necessarily experiential. But Kamm thinks that her explanation nevertheless holds considerable appeal. On her proposed way of looking at things, existing in a manner that is, on balance, neither good nor bad is preferable to posthumous non-existence, but the evil that accompanies a person’s posthumous non-existence need not be especially weighty. As she puts it: “It may be that *something*—even without any intrinsically positive features—is preferable to post-something *nothingness*, and yet, if something has intrinsically negative features (such as torture), these can override the negative of nothingness” (Kamm, 1993: 19, emphases added).

For many of us, Kamm’s view is attractive because it captures how we feel about the fact that there will come a day when we will no longer exist. Many people experience an abject terror when they contemplate how their death will put an end to their existence. While I have not experienced this terror myself, I understand that it is common, and very disturbing indeed.⁵

But despite its intuitive appeal, we should reject Kamm’s suggestion. For one thing,

⁵Philip Larkin describes the terror beautifully in his poem *Aubade* (Larkin, 1988: 208-9). As he puts it, when he lies awake at night, his mind sometimes blanks

“[...] at the total emptiness for ever,
The sure extinction that we travel to
And shall be lost in always. Not to be here,
Not to be anywhere,
And soon; nothing more terrible, nothing more true.”

if we take seriously her proposal that posthumous non-existence is intrinsically bad, then it is difficult to see how we can mitigate this badness by living longer, but still finite, lives. To see what I mean, suppose first that the proposed badness of our posthumous non-existence is unbounded—it keeps increasing without approaching some finite limit. If this is the case, then no matter how long our finite lives, and no matter how trivial the badness of nothingness that accrues over some finite amount of time, the badness of our posthumous non-existence will nevertheless remain infinite. Alternatively, suppose that the badness of our posthumous non-existence either approaches, or at some point reaches, a finite limit. If this is the case, then this finite limit will eventually be arbitrarily closely approached or reached, once again no matter how lengthy our finite lives, and no matter how trivial the badness of nothingness that accrues over some finite amount of time. Even if our posthumous existence is intrinsically bad for us, this would thus not seem to provide us with a reason in favour of choosing the longer life in *Two Lives*.⁶

In response, one might point out that our posthumous nothingness might be finite if the universe itself is finite.⁷ The thought here is that the intrinsic badness of a person's non-existence might continue to accrue from the day a person dies until the universe eventually ends. If this is the right way of looking at things, then living longer will, in

⁶One might try to escape this dilemma by suggesting that the alternative of living the longer life *time-slice dominates* the alternative of living the shorter life, and is for this reason preferable. An alternative *A* *time-slice dominates* an alternative *B* just in case *A* is at no times worse than *B*, and is at least during one time interval better than *B*. There are two problems with this suggestion. For one thing, there is no guarantee that living a longer life will at no time be worse than living a shorter life, even if we stipulate that the two lives in question have the same overall value, and even if we grant that it is intrinsically bad to be dead. Consider, for example, a very volatile long life that is at certain times very bad, but at other times very good. When such a life is bad, it might well be worse than non-existence, in which case it does not time-slice dominate an equally valuable shorter life. The argument thus succeeds at most for a proper subset of relevant cases. Second, if it is bad to be dead, then this badness is necessarily non-experiential, and it is not clear that such badness can meaningfully be said to accrue during clearly specified time slices, as the argument from time-slice dominance presupposes. For a defence of what Steven Luper (2016: sec. 4.6) calls *indefinitism* about the temporal location of non-experiential value, see e.g. Grey (1999). I thank Johanna Thoma for discussion on this point.

⁷I thank three reviewers of this journal for independently raising this point.

fact, reduce the amount of posthumous badness that a person is exposed to. Yet if Kamm (1993: 19) is correct and what is bad for a person is the “nothingness” that follows her existence, then it is not clear why this nothingness should depend for its reality on the existence of something else. Having said that, it is at least conceivable that the badness of the nothingness that Kamm has in mind necessarily has a temporal dimension, and that time ends once the universe collapses,⁸ in which case Kamm’s argument can be rendered valid. Yet despite this potential validity, I believe that we still have good reason to reject Kamm’s argument for the implausibility of its implications.

Consider the following. If posthumous non-existence were intrinsically bad, then this would provide us with a potentially very weighty reason against bringing people into existence. But while it is plausible that certain considerations can count against bringing people into existence, the posthumous nothingness that any person will inevitably be exposed to does not seem to figure among these considerations. Moreover, if posthumous non-existence were intrinsically bad, this would provide us with a reason in favour of postponing procreation, and it would entail that our ancestors were more unfortunate than we are, simply for having existed earlier, and for thus being exposed to more nothingness before the collapse of the universe.⁹ Neither of these implications seem plausible.

In principle, it might be possible to resist the just mentioned implications by claiming that non-existence is bad not just posthumously, but whenever it occurs. Having said that, as Kamm (1993: 20) herself points out, if non-existence is *bad for a person*, then we need to be able to designate a subject for whom it is bad not to exist. But if non-existence is bad even before someone is born, or even if someone is never born, then there might be no subject for whom this non-existence could be bad. Kamm thus distances herself from the claim that non-existence is bad whenever it occurs, and does so for good reason.

In sum, while Kamm’s proposal can potentially be rendered valid, we still have reason to reject it for the implausibility of its implications. We do better if we side

⁸According to Battersby (2011), there is no agreement among scientists whether time will end when the universe ends.

⁹I thank an anonymous reviewer for drawing out these two implications.

with proponents of the Deprivation Account who assume that, *pace* Kamm, the state of being dead lacks intrinsic value if we cease to exist when we die.

4 The Option Value of Life

4.1 Option Value in Investment Theory

Suppose that you are thinking about buying a laptop. You have found a model that suits your needs—it has a long battery life, it is lightweight, and it is fast. There is no cheaper model on the market that suits your needs equally well. Should you buy the laptop?

Buying the laptop would mean making an investment. To make an investment generally means to incur upfront costs in order to reap benefits later. According to the *discounted cash flow (DCF) method*—a method that formed the gold standard of investment theory from the 1970s to the mid-1990s, thus underlying the *neoclassical investment models* developed during that time—an investor should make an investment just in case (Bernanke, 1983: 90; Dixit and Pindyck, 1994: 4-5):

1. The investment's *net present value* is greater than zero.
2. The investment's net present value is no smaller than the net present value of other available investments.

An investment's net present value (NPV) is the *sum of all the expected costs and benefits*¹⁰ associated with the investment, expressed in a common currency, with future cash flows discounted to their *present value*.¹¹

¹⁰This analysis assumes that investors should be risk neutral. In this paper, I follow this assumption and presume that under uncertainty, a rational agent is risk neutral both with respect to financial value as well as with respect to 'life value'. I assume this only so as not to unnecessarily complicate the discussion in the main text; all the substantial points put forward are compatible with views according to which rational agents may exhibit attitudes towards risk that range from at least moderately risk-averse to at least moderately risk-seeking (see Dixit, 1992: 110).

¹¹Discounting future cash flows to their present value is important because there exists a market for capital. This market means that if you have money that you do not currently need to fund your

In the laptop example, if you want to know whether buying a laptop is a good investment, the DCF method tells you to list all the expected costs and benefits of your investment, convert them into a common currency, and adjust all amounts to their present value. Suppose you follow these instructions, and you come up with the following:

1. **Costs.** The upfront cost of the laptop is £1,000. You expect no further costs.
2. **Benefits.** You expect that your laptop will last you for three years. You expect that it will generate daily benefits that you value at £3 per day, expressed in present value. The total expected benefits are therefore $365 \cdot 3 \cdot £3 = £3,285$.

The total sum of the expected costs and benefits of buying the laptop, expressed in today's pound sterling, is $-\text{£}1,000 + \text{£}3,285 = \text{£}2,285$. This is clearly greater than zero. Assuming that there is no competing alternative investment that has an even higher NPV, you should buy the laptop.

The powerful idea behind neoclassical investment theory is that all the costs and benefits that are associated with an investment need to be converted into one and the same unit of measurement, so that they can be meaningfully compared. But while neoclassical investment theory captures a powerful idea, it has been criticised by advocates of the so-called *real options approach* to investment valuation for neglecting a set of considerations that are frequently crucially important. In this way, neoclassical investment theory can consumption or your investment projects, you can temporarily make this money available to others against a fee. When you consider the value of an investment, you need to take this opportunity cost into account. If others would pay you more to borrow money from you than you can expect to gain through investing it in a project of your own, then you are on expectation better off, other things equal, if you lend your funds to others. Especially if an investment is costly upfront, and promises significant financial benefits only much later, lending your money to others can be financially preferable.

In the context of this paper, I am interested in investments whose attractiveness is best assessed not in terms of their *monetary value*, but in terms of their *cardinally measurable life value* (see footnote 2). Following John Broome, I assume that no discount factor should be applied to future costs and benefits when costs and benefits are measured in terms of their impact on the value of a person's life (Broome, 2004: 68-76; see also Broome, 2009). As I explain in footnote 4, this is a controversial claim, but my argument does not depend on it.

sometimes lead us to invest too readily (see e.g. Henry, 1974a; Cukierman, 1980; Pindyck, 1991; Dixit, 1992).

To develop an understanding of the considerations that neoclassical investment models overlook, consider again the laptop example. Suppose you currently find yourself in the following situation. You are a finishing PhD student, and while you might stay in academia, you are also taking seriously a particular option to leave academia. More specifically, two days ago, you had a job interview with a small consulting firm whose work you respect. If they offer you a job, you will accept it. Otherwise, you will aim to stay in academia. When calculating the expected benefits of your laptop, you reasoned as follows: “If I get the job, I will not have any use for the laptop that I am considering buying, as I would be provided with a company laptop. Hence on the assumption that I get the job, which I would start a month from now, a new laptop would benefit me only for the next month. If I remain an academic, I estimate that the laptop will generate benefits for me for the next three years that I value at £5.84 per day. There is about a 50% chance that I will get the job. Hence the expected benefits of buying the laptop sum to an average of £3 per day.”¹² You will hear back from the consulting firm within the next four days.

Common sense demands that you wait until you hear back from the consulting firm before you decide whether to buy the laptop. But neoclassical investment theory cannot make sense of this fact. What is it missing out on? The real options approach to investment analysis draws our attention to the following key considerations (see e.g. Dixit and Pindyck, 1994: 6-9; Bloom, 2014: 163-5):

1. That you can *postpone your decision to invest*, and that the costs of postponement, i.e. the returns that you forego by not investing right away, may not be very high;
2. That your investment is at least partially *irreversible*. An investment is *fully irreversible* if, once you have decided to invest, you cannot withdraw from the investment, but have to let it run its course. An investment is *partially irreversible*

¹²The expected benefits are $30 \cdot £5.84 + (365 - 30 + 2 \cdot 365) \cdot 0.5 \cdot £5.84 = £3,285$, or roughly £3 per day for three years.

if you can withdraw from it, but doing so leaves you worse off than you would have been had you not invested;¹³

3. That there is *uncertainty about future costs and benefits* that will at least partially resolve itself within a specified time frame.

In combination, these three factors can give rise to a positive *option value* that attaches to not investing right away. If the cost of not investing right away is low, and if an investment is highly irreversible, this option value frequently means that you should not invest right away, even in situations where the NPV of an investment is positive. As R. Glenn Hubbard puts it, the real options approach to investment analysis formally develops the insight that sometimes one does best if one “keep[s] one’s options open” (Hubbard, 1994: 1816).

In the laptop example, the option value of waiting for four days before deciding whether to buy a laptop is calculated as follows. Four days from now, if you are made a job offer, you will decide *not* to buy the laptop. Hence there is a 0.5 probability that you will choose not to invest at all, in which case you will have to go without a new laptop for the remaining 26 days before you start your new job. This yields $-0.5 \cdot 26 \cdot \text{£}5.84 = \text{£}-75.92$. If no job offer is forthcoming, you will buy the laptop as soon as this is made clear to you. This yields $0.5 \cdot (-\text{£}1,000 + (365 - 4 + 365 \cdot 2) \cdot \text{£}5.84) \approx \text{£}2,686$. $\text{£}2,686 - \text{£}75.92 \approx \text{£}2,610$ is therefore the value, as seen from the point of view of today, of the option of being able to decide four days from now whether you should buy the laptop. This value is clearly positive, and it is larger than the value of buying the laptop without waiting, which is $\text{£}2,285$. It follows that you should decide to wait just in case postponing the decision is not too costly. In your case, waiting for four days will cost you $\text{£}5.84$ per day, or $\text{£}23.36$ in total. As $\text{£}2,610 - \text{£}23.36 \approx \text{£}2587$, the cost of waiting is low enough to make waiting the preferable strategy—the net value of waiting, i.e. its option value minus its cost, exceeds the value of buying right away by $(\text{£}2,610 - \text{£}23.36) - \text{£}2,285 \approx \text{£}302$.

¹³For the laptop example, I implicitly assume that your investment would be irreversible. The thought here is that if you buy the laptop, you would not later attempt to resell it, as the expected costs of reselling it exceed the expected benefits.

Early contributions to the real options literature focused on environmental issues.¹⁴ They proceeded from the insight that if an investment project involved some degree of environmental degradation, then this fact rendered the project at least partially irreversible. Based on this insight, they developed models to show that applying neoclassical investment theory to such projects risked leading to overinvestment, and to a corresponding under-protection of environmental resources (see e.g. Arrow and Fisher, 1974; Henry, 1974a; Henry, 1974b). Soon thereafter, the real options approach was extended to firm decisions (see e.g. Cukierman, 1980; Bernanke, 1983; McDonald and Siegel, 1986). In the 1990s, Avinash Dixit and Robert Pindyck played an important role in systematising, reviewing, and advancing the existing literature (see Pindyck, 1991; Dixit, 1992; Dixit and Pindyck, 1994). While they pointed out that the real options approach might hold important insights also for non-standard, non-economic ‘investment problems’ such as marriage and suicide, they limited themselves to offering a few brief qualitative speculations (see Dixit, 1992: 127; Dixit and Pindyck, 1994: 23-5). Since the 1990s, the real options literature has continued to grow at a fast rate (for recent review articles, see e.g. Bloom, 2014; Trigeorgis and Reuer, 2017). Different contributions to the literature model different decisional set-ups, or try to estimate empirically how significant a role the value of keeping one’s options open plays in different industries or contexts. While the detailed results of different contributions diverge due to varying modelling assumptions and empirical input, the following results and observations—all of which are important for our purposes—are fundamental and obtain generally (see e.g. Dixit and Pindyck, 1994: 6-9; Bloom, 2014: 163-5):

1. A positive option value of waiting arises only in situations where a decision to invest is not one-off, but can be postponed. If waiting is not an option, there can be no positive value associated with it.
2. Just as crucially, an investment needs to be at least partially irreversible for the option value of waiting to be positive. If it is possible to withdraw from an investment at any time and to subsequently be as well off as one would have been

¹⁴I thank an anonymous reviewer for pointing out this fact.

had one not made the investment in the first place, there is no potential value to delaying the decision to invest.

3. Third, and once again just as crucially, waiting has a positive option value only if (i) it is currently uncertain what the returns of the investment will turn out to be, and only if (ii) you expect relevant news that will remove at least some of this uncertainty, at which point you can (iii) reconsider whether to invest. The basic intuition here is simple. On the one hand, if you have access to all the relevant information right away, you can calculate your optimal strategy right away, and proceed accordingly. There is no point to reconsidering the problem again later. On the other hand, if you do not have access to all the relevant information right away, but there will be no later point in time at which you can reconsider the issue while relying on improved information, there is no point to delaying the decision either, as you will not be able to make an improved decision later.

In a highly influential paper, Ben Bernanke (1983: esp. 91-4) proves a very general ‘bad news principle’, which shows that the option value of waiting to invest depends crucially on the possibility of ‘bad news’, but is largely independent of the possibility of ‘good news’. *Bad news* are news that render an investment less attractive than it is from your current point of view. *Good news* are news that have the opposite effect—they render an investment more attractive than it is from your current point of view. In the laptop example, if you are offered a job four days from now, this constitutes bad news, as it decreases the value of owning a laptop compared to a situation where you do not yet know whether you will be offered the job. By contrast, not being offered the job would constitute good news. On the face of it, one might think that both the potential for bad news and the potential for good news influence the value of waiting to invest: the former by rendering a cautious attitude more appropriate, thus raising the option value of waiting; the latter by increasing the potential upsides to investing right away, thus decreasing the option value of waiting. Bernanke (1983) shows that this thinking is mistaken. The key mechanism behind his bad news principle is easily explained. In all investment

set-ups in which option value gets off the ground, it is assumed that an investor can *react to new information* if earlier they have chosen *not* to invest. It follows that if there are good news, the investor can react by choosing to invest. For this reason, the possibility of good news fails to speak in favour of investing right away over delaying the decision. By contrast, the possibility of bad news speaks in favour of delaying an investment whenever there is a possible future scenario “in which an early attachment to [the investment] would be regretted” (Bernanke, 1983: 92; see also Arrow, 1968; Henry, 1974b). As Bernanke (1983) shows, the option value of waiting increases both with the probability of bad news, as well as with how detrimental the effect of such news would be on the investment’s return (see also the model in the appendix of this paper).

4. In cases where neoclassical investment theory recommends investing right away because the NPV of an investment is positive, the real options approach recommends investing right away only in the subset of cases in which the NPV is larger than the *option value of waiting* minus the *costs of waiting*. Most real options models stipulate that if you decide to delay an investment, you thereby forego whatever would have been the investment’s return for that decision period (see e.g. Cukierman, 1980: 466 or Pindyck, 1991: 1116-7). In the laptop example, if you postpone your decision by four days, you lose out on four days of laptop-related benefits that I assumed you value at £5.84 each day. If there are such waiting-related costs, it can be rational to invest right away even in cases where the option value of waiting is positive and exceeds the NPV of an investment. More precisely, if the difference between the option value of waiting and the NPV of an investment is small, or if the costs of waiting are significant, it is possible that one does best by investing right away (see also the model in the appendix).

In the remainder of this paper, I first translate the key ideas behind the real options approach as I have just introduced them to the non-economic context of life and death (section 4.2). I then consider some relevant practical implications (section 5).

4.2 Option Value and Life and Death Decisions

In section 2, I have argued that we cannot use the insights of the Deprivation Account to make sense of a preference in favour of living the longer life in *Two Lives*. But *Two Lives* is a highly artificial case, in the sense that you are *assumed to know for certain* how valuable a life there is in store for you if you choose to go on living. In real life, while we sometimes have a good idea of what our future is likely to look like, we can hardly ever be certain about this. Consider again the introductory example. As a person who is diagnosed with an incurable and progressive condition, you would likely try to develop an understanding of what is expecting you by consulting statistical information about patients who have received a similar diagnosis in the past. Based on your research, it may then be reasonable to conclude that your future likely looks quite bleak, but that a largely symptom-free and enjoyable future nevertheless remains a possibility. The introductory example can thus be read as a case of *Two Lives Under Uncertainty*: if you die today, you will have lived a life with a certain amount of value. If you choose to go on living, it is unclear how valuable a life you will eventually have lived. But given the information that you have access to, it seems rather unlikely that staying alive will add to the value of your life. How is it rational to think about the value of staying alive in such a more realistic setup?

To sharpen the discussion, suppose that after doing some research, you conclude that it is 90% likely that your future will be marked by pain and suffering, but that there is also a 10% chance that your future will be largely symptom-free and enjoyable. Further suppose that you assign the following values (see footnote 2) to different possible lives: if you were to die today, you think that your life would have had an overall value of 500. It would clearly have been a life worth living, albeit a rather short one. If you do not die today and a painful future materialises, you estimate that the life that you will ultimately have lived will have an overall value of 400. You base this estimate on the thought that years of pain and suffering would detract from the current value of your life. Finally, if you do not die today and your future turns out to be enjoyable, you estimate that the life that you will ultimately have lived will have an overall value of 600. With the case

thus sharpened, should you conclude that it is in your self-interest to go on living?

Taken by itself, the Deprivation Account cannot be used to answer this question. According to the Deprivation Account, it is in your self-interest to go on living just in case you will eventually have lived a better life if you do so. It is thus prudentially rational to go on living just in case your future will be enjoyable. But the crux is that you do not know what lies in your future. To generate an action-guiding verdict, we thus have to combine the insights of the Deprivation Account with a suitable rule for making decisions under uncertainty.

Suppose first that we combine the key ideas behind the Deprivation Account with the core insights of neoclassical investment theory. On this way of looking at things, choosing to die means *choosing to disinvest from life*. Life represents an ongoing investment that keeps generating positive or negative returns as you keep accruing positive or negative life value.¹⁵ If the negative life value that accrues outweighs the positive life value, the project of staying alive operates at a loss, and discontinuing one's investment may be reasonable.

According to neoclassical investment theory, you should choose to disinvest from life by committing suicide just in case the net present value (NPV) of suicide, i.e. the sum of all of its associated expected costs and benefits, is greater than zero (see section 4.1 above). In the case at hand, the NPV of choosing to die is the probability-weighted negative life value that you avoid if a painful future awaits you, minus the probability-weighted positive life value that you miss out on if an enjoyable future awaits you, i.e. $0.9 \cdot 100 - 0.1 \cdot 100 = 80$. If there are any costs associated with committing suicide that affect the value of your life negatively, they represent an upfront cost of the choice to disinvest from life that reduce its NPV. If, for example, committing suicide reduces the value of your life by 5 “life value” points because it is a shameful way to die, then the NPV of choosing to die is reduced to $-5 + 0.9 \cdot 100 - 0.1 \cdot 100 = 75$. More generally,

¹⁵Alternatively, we can think of death itself as an investment that keeps generating positive or negative returns, where the positive returns are negative life value avoided, and the negative returns are positive life value missed out on. What terminology we apply is inconsequential (see e.g. McDonald and Siegel, 1986: 711).

in the case at hand, neoclassical investment theory yields the verdict that you should kill yourself just in case the upfront cost of suicide lies below 80 life value points.

This way of thinking about the rationality of suicide is rather widespread, even if the terminology of neoclassical investment theory is not always introduced. Hamermesh and Soss (1974) develop an explicitly economic framework to argue that it is rational to end one's life if the expected value of the rest of one's life is negative. Narveson (1986) agrees, though he chooses a non-formal approach to make the same point. Frequently, the assumption that it is rational to end one's life if one does not expect to benefit from continuing to live forms an implicit backdrop to discussions about suicide, though it is neither explicitly mentioned nor defended (see Singer, 2003, for a relevant example). The problem with this way of looking at things is that it can lead us to 'overinvest' in suicide by making suicide look more attractive than it really is (Dixit and Pindyck, 1994: 24). More specifically, when we can choose between suicide and continuing to live, this choice situation may have all of the features that are needed to give rise to a positive option value of not ending one's life right away. When this is the case, staying alive may be in a person's self-interest even if the NPV of dying right away is greater than zero.

To see why a choice between life and death may have all of the features that are necessary to give rise to a positive option value, consider again the situation that you would find yourself in if diagnosed with a progressive and incurable disease. First, if you choose to die, you choose something irreversible. Once you disinvest from life, you cannot subsequently choose to reinvest. Instead, you are stuck with whatever cost-benefit profile your disinvestment generates, i.e. with whatever difference there exists in life value between the life that you actually lived, and the life that you would have lived had you chosen to remain alive. Next, if you choose to remain alive, you may well be able to reconsider this decision at later points in time. Finally, if you choose to remain alive, it seems likely that you will eventually find out whether your symptoms turn out to be mild or pronounced. If all of these conditions are in place, then there will be a positive option value to staying alive, as killing yourself right away is a regrettable choice in the unlikely but far from impossible scenario in which your symptoms would only ever have

been very mild (Bernanke, 1983).

When there is a positive option value to staying alive, you should choose to stay alive just in case the option value of staying alive minus the costs of staying alive exceed the NPV of killing yourself right away. There are costs to staying alive if you expect that you will accrue negative net life value between your decision not to kill yourself right away and the time in the future at which you plan to reconsider this decision. Suppose that in order to find out whether your disease-related symptoms can be managed, you will have to go through a painful period where your physicians try to fine-tune combinations of different medications so as to fit your idiosyncratic needs. If such a painful calibration period awaits you, then there are significant costs to delaying your decision to disinvest from life. By contrast, if you can simply wait and see to find out what the future has in store for you, the costs of postponing your decision are likely negligible.

In the appendix, I present a simple toy model for the option value of life that is inspired by Bernanke (1983) and Dixit and Pindyck (1994). The model lays out the key dimensions that give rise to, and help determine the size of, the option value of staying alive.

Once we move beyond the simple set-up of the *Two Lives Under Uncertainty* case in which you are diagnosed with an incurable and progressive disease, the following observation is noteworthy. In situations where you are contemplating whether killing yourself might be rational, the rationality of suicide depends crucially on the particular reasons why you think that continuing to live might not add to the value of your life. As clarified in section 4.1 above, a positive option value to staying alive arises only in cases where there is (i) uncertainty about the value of continuing to live that you expect will (ii) resolve itself at a point in time at which you can (iii) reconsider your decision to stay alive. More precisely, as Bernanke's bad news principle clarifies, there is positive option value to staying alive if it is possible that killing yourself right away might turn out to be a regrettable choice, and if you will be able to control potential downsides to staying alive by killing yourself if and when it becomes clear that such downsides will in fact materialise. Whether this is the case would depend crucially on the crisis you are facing. In *Two*

Lives Under Uncertainty, the crisis you are facing demands that you come to terms with the fact that your future will most likely be very painful. There is a potential upside to staying alive—you might not be hit by pain—and it seems reasonable to assume that you can manage the potential downsides of staying alive by reconsidering your decision to stay alive at suitable future points in time. If you are not in any immediate and unbearable pain, it is thus most likely not in your self-interest to kill yourself right away. By contrast, suppose that you have fallen in love with your best friend's partner, and you are confident that your feelings are reciprocated. You think that if you stay alive, it is highly likely that you will betray your best friend. At the same time, you think that betraying your best friend would irredeemably blight your life. In this admittedly quite dramatic scenario, it is much less clear that you should choose to remain alive, at least if your ideas about what makes a person's life go well or badly are not mistaken. While there is a potential upside to staying alive—you might not betray your friend, and your love interest might eventually change—it is doubtful that the continued possibility of suicide provides an effective tool for keeping in check the potential downsides of staying alive. After all, you fear that if you choose to stay alive, your lack of self-control will almost inevitably lead you to act in a way that you think is unforgivable. While the option value of life might thus frequently play an important role in crises that are related to one's physical health, its significance in the context of other crises might be more limited.

5 Practical Implications of the Option Value of Life

It might appear that much of the discussion in this paper is of limited practical significance. After all, people who are feeling suicidal rarely consider in a calm and reasoned manner whether it would be in their self-interest to commit suicide. They rather feel very strongly that suicide is their only viable option under the circumstances, and the idea that there might be an option value to staying alive is unlikely to appear to them forceful and pertinent. Similarly, as long as we are not feeling suicidal, we tend to feel very attached to our lives, and we do not need to be told that we might have good reason for wanting to

stay alive, even if we have just been diagnosed with a most likely incurable and progressive condition. Apart from the rare occurrences where a rather unusual individual tries to establish in a careful manner whether staying alive is in her self-interest, the option value of life might thus appear fairly inconsequential. But this way of looking at things misses a number of important considerations.

For one thing, a proper understanding of the option value of life might help justify targeted paternalistic interventions aimed at preventing irrational, but not rational, suicides. In their brief discussion of suicide, Dixit and Pindyck (1994: 24-5) suggest that it might be a good thing if, as a society, we condemn suicide as highly immoral. As they put it:

“Suicides project the bleak present into an equally bleak future, ignoring uncertainty, and thereby ignoring the option value of life. Then religious or social proscriptions against suicide serve a useful function as measures to compensate for this failure of rationality. By raising the perceived cost of the act, these taboos lower the threshold of quality of life that leads to suicide when option value is ignored. This can correct the failure of the individuals’ forethought, and bring their threshold in conformity with the optimal rule that recognizes the option value” (Dixit and Pindyck, 1994: 24-5).

Pace Dixit and Pindyck, we should not be quick to embrace religious and social proscriptions against suicide. Equipped with an understanding of option value, we can appreciate that suicides which “project [a] bleak present into an equally bleak future, ignoring uncertainty,” might frequently be irrational, as Dixit and Pindyck point out. But we can appreciate as well that the choice to die can be rational. More precisely, if a person is not currently living a life worth living—maybe because they are experiencing a crushing sense of despair or anguish—this represents a cost to delaying the decision to die further into the future. If this cost is large enough, this can render suicide rational, especially if there is only a rather slim chance that things will improve substantially. By speaking forcefully against irrational *and rational* suicides alike, and by not encouraging us to try to distinguish between them, societal taboos against suicide risk increasing or

prolonging unnecessary suffering in cases where suicide is rationally advisable. Even if Dixit and Pindyck are right that suicidal individuals tend to exaggerate the extent to which their situation is hopeless, our best reaction to this might thus not be to make suicide very costly. Instead, we might fare better if we aim to mitigate the ongoing suffering that suicidal individuals experience, and if we help them investigate why their situations might not be as dire as they seem.

The fact that there can be an option value to life has practical implications not only in the context of suicide, but also with respect to voluntary and non-voluntary euthanasia, where it provides a reason in favour of increasing their availability.¹⁶ To see this more clearly, suppose that one of your family members is in hospital. They are not currently in a responsive state. Maybe they were involved in an accident, and are now in a coma. The doctor calls you and your other family members into her office. According to the doctor, the outlook for the patient is bleak. There is a good chance that the patient will not wake up from their coma, and if they do wake up, their brain will likely be severely damaged, and their life not worth living. You ask the doctor whether there is a chance that the patient will wake up and have a good life. The doctor says that this is not impossible, but that the chances are slim. The doctor then asks you: “Should we introduce life-sustaining measures?”

With respect to a situation such as this one, focusing on option value brings out two things. First, it brings out that there is an important respect in which it is valuable if voluntary and non-voluntary euthanasia are available to a patient if the patient wakes up into a life that is not worth living (see Savulescu, 1994). If such interventions are not available, then the earlier decision to keep the patient alive cannot be revisited in a scenario in which the patient wakes up into a life not worth living. This means that the downward risks of keeping the patient alive cannot be controlled. This, in turn, may well render it in the patient’s best interest not to introduce any life-sustaining measures, or to disinvest from life right away. By contrast, if voluntary and non-voluntary euthanasia

¹⁶If Dixit and Pindyck (1994: 24-5) are right that many suicides are based on the irrational projection of a bleak present into an equally bleak future, it is crucial that voluntary euthanasia should be made available only in cases where it has been established that an individual’s wish to die does not rest on such an irrational projection. I thank an anonymous reviewer for drawing my attention to this point.

are available should the patient wake up into a life not worth living, there most likely is an option value to staying alive, and it most likely is in the patient's best interest to remain alive. The availability of euthanasia therefore increases the attractiveness of the patient's available alternatives.

Second, focusing on option value brings out that if euthanasia is available, then the patient's family are not merely clutching at straws, or unable to let go, if they insist on life-sustaining measures. Instead, as long as the patient is not in pain, it is in the patient's best interest to remain alive, and the family are safeguarding the patient's interests by keeping him or her around.¹⁷ A family who decides that life-sustaining measures should be introduced is thus not condemning their loved one to the pitiable existence of a 'human vegetable' because they find themselves unable to let go. If we take the idea of option value seriously, and if we let ourselves be inspired by Emily Dickinson's poem, we can instead say that a family who votes in favour of life-sustaining measures are thereby enabling their loved one to "dwell in possibility".

ACKNOWLEDGEMENTS

For helpful written comments and discussions on earlier versions of this paper, I thank Ben Ferguson, Roberto Fumagalli, Margherita Harris, Todd Karhu, Johanna Thoma, Thomas Seiler, and three anonymous referees of this paper. I owe special thanks to Bryan Roberts and Thomas Seiler for providing valuable help with the model in the appendix of this paper. Earlier versions of the paper were presented at the *Oxford Moral Philosophy Seminar* and at the *8th Braga Meetings on Ethics and Political Philosophy*. I thank the respective audiences for their questions and comments.

¹⁷Even if keeping the patient alive is in the *patient's* interest, it need not be the right thing to do *all things considered*. If it is very costly to keep the patient alive, the relevant resources might be better spent elsewhere, even if such a conclusion is very difficult to accept.

References

- Arrow, K. J. 1968. Optimal Capital Policy and Irreversible Investment. In Wolfe, J. N., editor, *Value, Capital, and Growth*, pages 1–20. Aldine, Chicago.
- Arrow, K. J. and Fisher, A. C. 1974. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics*, 88(2): 312–19.
- Battersby, S. 2011. Will time end? *New Scientist*, (898, 8 October issue).
- Bernanke, B. S. 1983. Irreversibility, uncertainty, and cyclical investment. *Quarterly Journal of Economics*, 98(1): 85–106.
- Bloom, N. 2014. Fluctuations in uncertainty. *Journal of Economic Perspectives*, 28: 153–176.
- Bradley, B. 2009. *Well-Being and Death*. Oxford University Press, Oxford.
- Broome, J. 2004. *Weighing Lives*. Oxford University Press, Oxford.
- Broome, J. 2009. Discounting the Future. In *Ethics out of Economics*, pages 44–67. Cambridge University Press.
- Cukierman, A. 1980. The effects of uncertainty on investment under risk neutrality with endogenous information. *Journal of Political Economy*, 88(3): 462–75.
- Dixit, A. 1992. Investment and hysteresis. *Journal of Economic Perspectives*, 6: 107–132.
- Dixit, A. and Pindyck, R. 1994. *Investment Under Uncertainty*. Princeton University Press.
- Feldman, F. 1992. *Confrontations With the Reaper. A Philosophical Study of the Nature and Value of Death*. Oxford University Press, New York NY.
- Grey, W. 1999. Epicurus and the harm of death. *Australasian Journal of Philosophy*, 77(3):358–64.

- Hamermesh, D. S. and Soss, N. M. 1974. An economic theory of suicide. *Journal of Political Economy*, 82: 83–90.
- Henry, C. 1974a. Investment decisions under uncertainty: the “irreversibility effect”. *American Economic Review*, 64(6): 1006–12.
- Henry, C. 1974b. Option value in the economics of irreplaceable assets. *Review of Economic Studies Symposium*, 41: 89–104.
- Hubbard, R. G. 1994. Investment under uncertainty: keeping one’s options open. *Journal of Economic Literature*, 32(4):1816–1831.
- Kamm, F. 1993. *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford University Press, New York NY.
- Kamm, F. M. 2017. The purpose of my death: death, dying, and meaning. *Ethics*, 127:733–761.
- Larkin, P. 1988. *Collected Poems*. Faber & Faber, London.
- Luper, S. 2016. Death. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2016 edition.
- McDonald, R. and Siegel, D. 1986. The value of waiting to invest. *Quarterly Journal of Economics*, 101(4): 707–28.
- McMahan, J. 2002. *The Ethics of Killing. Problems at the Margins of Life*. Oxford University Press, New York NY.
- McMahan, J. 2019a. Early Death and Later Suffering. In Gamlund, E. and Solberg, C. T., editors, *Saving People from the Harm of Death*, pages 116–133. Oxford University Press, Oxford.
- McMahan, J. 2019b. In Defense of the Time-Relative Interest Account. In Gamlund, E. and Solberg, C. T., editors, *Saving People from the Harm of Death*, pages 267–278. Oxford University Press, Oxford.

Nagel, T. 1970. Death. *Noûs*, 4: 73–80.

Narveson, J. 1986. Moral philosophy and suicide. *Canadian Journal of Psychiatry*, 31(2): 104–7.

Pindyck, R. S. 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, 29: 1110–1148.

Savulescu, J. 1994. Treatment limitation decisions under uncertainty: the value of subsequent euthanasia. *Bioethics*, 8(1):49–73.

Singer, P. 2003. Voluntary euthanasia: a utilitarian perspective. *Bioethics*, 17:526–41.

Trigeorgis, L. and Reuer, J. J. 2017. Real options theory in strategic management. *Strategic Management Journal*, 38: 42–63.

Appendix

A toy model for the option value of life can be drawn out of Bernanke’s (1983) ‘bad news principle’ model of irreversible investment, following the simplifications introduced by Dixit and Pindyck (1994: 40-1). The model provides a decision procedure for choosing to disinvest from life by investing in death through suicide or euthanasia. In a life-and-death context, Bernanke’s ‘bad news principle’ becomes a ‘good news principle’: the decision to commit suicide depends only on the amount of good news that one anticipates for the future, but not on the amount of bad news. For simplicity, we consider just two time intervals, I_0 and I_1 . The first time interval will be interpreted as a period at the beginning of which you choose between death and continued life. The second is a later time period at the beginning of which you receive some clarification about the remaining value of your life, and where you can make the same choice again, if you are still alive.

The Model

Let $R_0 < 0$ be the value that accrues if one is alive during time interval I_0 . R_0 is assumed to be negative because the badness that accrues during I_0 is assumed to outweigh the goodness that accrues during I_0 , for example because I_0 represents a period characterized by intense suffering. $-R_0$ represents the value of being dead during I_0 , as one avoids being exposed to R_0 by being dead. Let R_1 be a random variable representing the value that accrues if one is alive during I_1 , which we assume satisfies,

$$R_1 = \begin{cases} bR_0 & \text{with probability } p, \\ gR_0 & \text{with probability } 1 - p. \end{cases} \quad (1)$$

for some $b > 0$, $g \leq 0$ ('bad news' and 'good news' parameters, respectively), and for some $0 \leq p \leq 1$.

We begin by characterising the net present value (NPV) of choosing death at the beginning of I_0 . Ignoring discounting (see footnote 11), the NPV is given by adding together the value of being dead during I_0 , i.e. $-R_0$, with the expected value of being dead during I_1 , which is given by the probability-weighted values of avoiding the good-news and bad-news scenarios (Equation 2):

$$\begin{aligned} \text{NPV}_d &= -R_0 - p \underbrace{(bR_0)}_{\text{bad news}} - (1 - p) \underbrace{(gR_0)}_{\text{good news}} \\ &= (-1 - g + p(g - b))R_0. \end{aligned} \quad (2)$$

Next, to characterise the net present value of choosing life at the beginning of I_0 , and thus facing the choice again at the beginning of I_1 , we must sum contributions from both the bad news and good news scenarios.

First scenario: suppose we choose life at I_0 , and this is followed by bad news (which occurs with probability p). Then at I_1 , we can either choose to live (and accrue value of bR_0 , as given by Equation 1), or we can choose to die (and avoid the same value). We assume a strategy in which the option of greater value is always chosen, and so the

resulting net value is given by: $\max[bR_0, -bR_0]$. As $b > 0$, $R_0 < 0$, death will be chosen at I_1 under the bad news scenario.

Second scenario: suppose we choose life at I_0 , and that this is followed by good news (which occurs with probability $(1 - p)$). Then at I_1 , a similar argument results in a net value of: $\max[gR_0, -gR_0]$. As $g \leq 0$, $R_0 < 0$, life will (weakly) be chosen at I_1 under the good news scenario.

The total NPV for choosing life at I_0 is now given by the probability-weighted sum of these two scenarios, plus the value accrued during I_0 :

$$\text{NPV}_l = R_0 + p(\underbrace{-bR_0}_{\text{bad news}}) + (1 - p) \underbrace{gR_0}_{\text{good news}} \quad (3)$$

Now let us derive the condition under which $\text{NPV}_d = \text{NPV}_l$. Under this condition one would be indifferent between death and staying alive at I_0 .

After simplifying, equating Equations 2 and 3 gives

$$(p - 1)g = 1. \quad (4)$$

If the value of $(p - 1)g < 1$, then it is preferable to die, since then $\text{NPV}_d > \text{NPV}_l$. If the value of $(p - 1)g > 1$, then it is preferable to live, since then $\text{NPV}_l > \text{NPV}_d$. If $(p - 1)g = 1$, a rational agent is indifferent between dying right away and staying alive.

Properties of the model

This model has at least the following interesting properties, which include a Bernanke-style ‘good-news principle’.

Departure from neoclassical investment theory

Suppose that $\text{NPV}_d > 0$. In line with neoclassical investment theory, one might take this to be reason enough to commit suicide at the present moment (see section 4). However, this is not a prediction of the model. The model rather advises that one should often wait

and continue to live to settle the possibility of future good news, which is not a strategy that neoclassical investment theory considers. More precisely, as long as $(p - 1)g > 1$, then no matter how large NPV_d , it is always more sensible to wait (as $NPV_d < NPV_l$).

Two-parameter decision procedure

On this model, in order to decide whether to commit suicide, one only needs to check whether $(p - 1)g < 1$. If yes, then it is preferable to die right away; otherwise it is weakly preferable to stay alive. This means that only the following parameters are needed:

- g : the amount of good news anticipated. More precisely, g indicates how the amount of net positive value that is expected to accrue during I_1 if good news materialise compares to the net negative value that accrues during I_0 .
- $1 - p$: the probability that good news will materialise.

Good news principle

Notably, the decision to commit suicide is independent of the amount of bad news that is expected. This analysis thus contains a ‘good news principle’: whether or not one chooses to die should not depend on the quantity of bad news anticipated, but only on the quantity of good news. This asymmetry arises out of an asymmetry of reversibility: whereas choosing to live is reversible, choosing to die is irreversible. This means that if one chooses to stay alive at I_0 , one can still take advantage of good news during I_1 , or choose to avoid the impact of bad news at the beginning of I_1 . By contrast, if one chooses suicide in I_0 , there is no such flexibility at I_1 .

Idealisations

The model makes the following idealisations, most of which can be easily removed:

- *Binary time*. In practice, a life has multiple decision times, not just two. However, this simplifying assumption can be easily removed and replaced with continuous

time models of irreversible investment, making the obvious adjustments we have made here for disinvestment; see e.g. Dixit and Pindyck (1994: ch. 5).

- *No discounting.* Future value might have to be discounted (see footnotes 4 and 11). Adjusting the model accordingly is simple, and does not change the central results; see Dixit and Pindyck (1994: 40-41).
- *Costless suicide.* It is assumed that suicide is costless. As with a discount factor, the model can easily be extended to accommodate such costs; see Dixit and Pindyck (1994: 40-41).
- *Time-independent variables.* It is assumed that the variables representing the amount of good news and bad news are time-independent. This assumption can be replaced in more complex models.
- *Cardinal and temporal assumptions about life value.* The model assumes that factors that make a person's life go well or badly can be quantified in a cardinal manner (see footnote 2) and accrue during identifiable time slices of a person's life (cf. footnote 6). It is moreover assumed that if the net value that accrues during a time interval I is ≥ 0 , then life is weakly worth living during that interval, but not otherwise. A formal analysis of the option value of life seems limited to theories about what makes a person's life go well that are consistent with this set of assumptions.

BIOGRAPHICAL INFORMATION

Susanne Burri is an Assistant Professor in the Department of Philosophy, Logic & Scientific Method at the London School of Economics and Political Science (LSE). Her main areas of research include normative ethics, decision-making under uncertainty, and the philosophy of death. Susanne Burri holds a B.A. in Economics from the University of St. Gallen in Switzerland (2009), an MSc in Economics and Philosophy from the LSE (2010), and a PhD in Philosophy from the LSE (2014).