Geometric Averaging in Consequentialist Ethics

Alfred Harwood

Abstract

When faced with uncertainty, consequentialists often advocate choosing the option with the largest expected utility, as calculated using the arithmetic average. I provide some arguments to suggest that instead, one should consider choosing the option with the largest geometric average of utility. I explore the difference between these two approaches in a variety of ethical dilemmas and argue that geometric averaging has some appealing properties as a normative decision-making tool.

1 Introduction

This paper is concerned with decision-making under uncertainty. The world is often uncertain and in order to make decisions, one must have a method for deciding between opportunities that do not have a guaranteed outcome. One way to do this is to use a mathematical process which averages over the uncertainty, allowing uncertain options to be represented by a single number. The option with the highest average is then chosen.

Commonly, in economics and decision theory, the arithmetic average, or *expected value* of a lottery is calculated and this value is used to quantify its desirability. If a lottery offers several possible monetary payouts x_i, each with respective probabilities p_i, the expected value of the lottery is [1]:

$$\mathbb{E}[x] = \sum_i p_i x_i \,. \tag{1}$$

In this paper, we will use the terms *arithmetic average*, *expected value*, and (for reasons that will shortly be explained) *ensemble average* to all refer to averages which take this form. We will compare this form of average to another form of averaging, known as *geometric averaging* [2]. We will denote the geometric average of the same lottery as:

$$\mathbb{G}[x] = \prod_i x_i^{p_i} \,. \tag{2}$$

[1]In this paper, we will restrict our consideration to discrete gambles, with a finite number of outcomes. Arithmetic averaging can be easily extended to cover continuous gambles by integrating over continuous probability distributions using a standard definite integral. Something similar can be done for geometric averaging by using the 'product integral' or 'geometric integral'. These are not covered here as they add mathematical complexity while adding no extra conceptual clarity.

[2]The quantities described in equations (1) and (2) are more accurately described as arithmetic and geometric *weighted means*. For compactness, we will refer to them as arithmetic and geometric *averages*.

While it is not as ubiquitous as the expected value, the geometric average as a decision-making tool has a long history going back at least to Bernouilli's 1738 paper[3]. In this paper, we will compare geometric and arithmetic averaging with respect to their qualities as normative rules of ethical decision-making.

Though it is often convenient to describe lotteries in terms of money, they do not have to concern wealth. When considering applied ethics, one often encounters uncertain situations where the outcomes are not specified in terms of money, but some other quantity that we care about. Instead of comparing gambles where the amount of money we win is uncertain, we may find ourselves comparing situations where a different, morally relevant quantity, (such as the number of lives we might save in a crisis situations) is uncertain.

Consequentialism is the name for a group of normative ethical views which hold that the goodness of an action should be judged only by its consequences. Often, the term *utility* is used to refer to a quantity of 'goodness' in a particular state of the world and consequentialism is framed as claiming that one's normative ethical duty is to choose actions which maximize the utility of the world. In this paper, we remain agnostic as to the precise definition of 'goodness' but use the examples used in this paper which are hopefully uncontroversial as they rely on concepts of goodness such as 'saving ten lives is, other things being equal, better than saving just one life'.

When presented with gambles with uncertain outcomes, it is often taken for granted that, if one is a consquentialist, one is morally obliged to choose the option with the highest 'expected utility' as calculated using the arithmetic average[4] [5]. In this way, the desirability of a probability distribution over many possible outcomes is boiled down to a single number and it is assumed that the arithmetic average is the best number for this purpose. It is assumed that the arithmetic average provides a normative standard by which one ought to make decisions. Here, we reject this assumption and explore the consequences of doing so. While there are several compelling reasons to favour the arithmetic average, we argue that these reasons are not overwhelmingly compelling when compared to the reasons for using geometric averaging.

In this paper, we ask two questions. Firstly, 'are there reasons to use geometric averaging, as opposed to arithmetic averaging, when making ethical decisions under uncertainty?' and secondly 'what kind of ethical behaviour would result from maximizing the geometric average of morally relevant quantities, as opposed to the arithmetic average?'.

Note that the averaging methods described here are by no means the only two in existence. The arithmetic average is discussed here as it is commonly used, and the geometric average has been chosen as a comparison due to its history, interesting properties, and relative simplicity. This study leaves untouched the rich vein of other averaging methods which exists in mathematics.

[3] Daniel Bernoulli. "Exposition of a New Theory on the Measurement of Risk". In: *The Kelly Capital Growth Investment Criterion: Theory and Practice.* World Scientific, 2011, pp. 11–24.

[4] R. A. Briggs. "Normative Theories of Rational Choice: Expected Utility". In: *The Stanford Encyclopedia of Philosophy.* Ed. by Edward N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University, 2019.

[5] See also Parfit (Derek Parfit. *Reasons and Persons.* OUP Oxford, 1984): 'What we ought subjectively to do is the act whose outcome has the greatest *expected* goodness'. Parfit then clarifies that, by expected goodness, he means the arithmetic average.

We make two notes before the study begins in earnest. Firstly, this paper takes a 'naive' view of probabilities. We assume that, for any lottery the probabilities are known in advance and these probabilities represent the 'true' proportion of occasions each outcome will occur, thus sidestepping the subtleties of the philosophical basis of probability. While such issues are fascinating, they affect arithmetic and geometric averaging equally and are thus not relevant to our discussion. Secondly, while this paper uses the language of consequentialism its results are not only of relevance to consequentialists. Non-consequentialists can and will find themselves in situations where they have to make decisions under uncertainty. A deontologist public health official may find herself having to decide between different policies for saving lives which each have uncertain outcomes. The virtue-ethicist captain of a sinking ship may find himself deciding how to save as many passengers as possible, in a situation where no outcomes are guaranteed.

2 Arithmetic Averaging

Before discussing the geometric average we will consider some reasons in favor of using the arithmetic average (also known as the expected value) as guide to decision-making under uncertainty. Let us introduce the *Arithmetic Averaging Decision Criterion* (AADC):

> **Arithmetic Averaging Decision Criterion.** When facing a decision between uncertain options, choose the option which has the highest arithmetic average of the quantity about which you care.

We will now summarize four arguments in favor of using the AADC for ethical decision-making.

2.1 The Von Neumann-Morgenstern Theorem

The Von Neumann–Morgenstern (VNM) utility theorem states that an agent whose behaviour satisfies four axioms of 'rational behaviour' will, when making decisions under uncertainty, behave as if she is maximising the expected value of a utility function, which assigns to each state of the world a number between zero and one[6]. To discuss these axioms further, we will introduce some notation which we will continue to use throughout. For lotteries A, B, C etc. let $A \prec B$ to mean 'lottery B is preferable to lottery A' and $A \sim B$ to indicate indifference between A and B. Furthermore let ApB indicate a mixture of lotteries A and B; a situation where lottery A occurs with probability p and lottery B occurs with probability $(1-p)$. Briefly, the axioms of the VNM theorem are *Completeness* (for two lotteries A and B either $A \prec B$ or $B \prec A$ or $A \sim B$), *Transitivity* (if $A \succ B$ and $B \succ C$, then $A \succ C$), *Independence* ($A \succ B$ if and only if $ApC \succ BpC$) and *Continuity* (if $A \succ B \succ C$, there exists p such that $ApC \sim B$). Each of these axioms of rational behaviour themselves can be justified through so-called 'money-pump' arguments[7] which demonstrate predictable ways in which agents can be exploited if they do not follow the VNM axioms. While the

[6] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, 2007.

[7] Johan E Gustafsson. *Money-pump Arguments.* Cambridge University Press, 2022.

VNM theorem does not concern ethics *per se* it does concern rational behaviour. If a person followed a normative theory that allowed her to be exploited into accepting states with lower utility such a theory would be self-defeating and hard to take seriously.

2.2 Ensemble Averaging

Taking the expected value of an uncertain quantity shares similarities with the process of 'ensemble averaging', used in statistical physics. In ensemble averaging, the average properties of a system are calculated by assuming a large, hypothetical 'ensemble' of identical systems, each in a possible state that could represent the 'true' state of the system. Average properties of the system are derived by averaging over this ensemble, with each member of the ensemble weighted by the likelihood that the system is in that particular state.

A similar approach can be applied to gambles, which we will demonstrate here with a simple example. Consider a gamble where a fair coin is tossed. If it lands heads, you will receive $100, and if it lands tails, you receive $0. The expected value of this gamble is $\$100 \times 0.5 + \$0 \times 0.5 = \$50$. This 'expected value' does not represent the value that one expects to receive, as winning $50 is not a possible outcome. Instead, it could be thought of as the profit you would receive if you teamed up with a large ensemble of other people (or hypothetical versions of yourself) who each independently took the gamble and agreed to share all profits evenly. As the size of the ensemble approaches infinity, the amount of money you will receive after it is shared will approach the expected value of the gamble. A key assumption in this case is that the probability of heads (50%) represents the proportion of members of the ensemble for whom the coin lands heads. The laws of large numbers can be invoked to argue that the larger the ensemble, the more likely this assumption is to be true. Note that, even if the practical conditions of a gamble prevent you from harnessing a large ensemble of other people, thinking about what would happen in such a situation can be used to give an idea of the general tendencies of the gamble.

2.3 Time Averaging with Additive Dynamics

In the previous section, we encountered the 'ensemble average' justification for using the expected value, using the term from statistical physics. In this section, we will look at a different kind of averaging method also used in physics–the time average. Time averaging does not require the use of a hypothetical ensemble of worlds and instead averages the behaviour of a single system over a long period of time.

Systems (or more precisely, observables of a system) for which the time average is equal to the ensemble average are described as 'ergodic'. For ergodic observables the time averages and ensemble averages can be used interchangeably. Broadly, this means that the ensemble-averaged properties of a system can be used to predict the properties it will converge to at equilibrium, after it is left for a long time. This cannot be done for non-ergodic systems. Recently, this distinction has been taken from physics and applied to economics and decision theory under the name 'ergodicity economics'[8] [9].

[8] Ole Peters. "The Ergodicity Problem in Economics". In: *Nature Physics* 15.12 (2019), pp. 1216–1221.

[9] In this study we set aside the question as to whether 'ergodicity economics' provides a good *descriptive* model of human behaviour (a view which is criticized in eg. (Jason N Doctor, Peter P Wakker, and Tong V

2.3.1 Additive and Multiplicative Dynamics

In order to find time-averages for a variable relating to a particular gamble, we must investigate the behaviour of this variable in the hypothetical situation where the gamble is repeated infinitely many times. To do this, we must specify the 'dynamics' of the gamble; that is, the relationship between the outcome of one gamble and the stake of the next. We will consider two intuitive types of dynamics: additive dynamics (where the *quantities* at stake in a gamble are kept constant in each repetition) and multiplicative dynamics (where the *proportions* at stake in each gamble are kept constant, relative to one's wealth *at the beginning of each repetition*).

As an example of these two types of dynamics, consider our simple gamble from earlier: a coin is tossed, and you receive \$100 if it lands heads and \$0 if it lands tails. Under additive repetition, each time that the gamble is repeated, you will receive \$100 when the coin lands heads and \$0 when it lands tails. For multiplicative dynamics, we need to know what proportion of your total wealth is represented by \$100. Under multiplicative repetition, the amount you would receive upon the coin landing heads in each repetition would be this proportion of your wealth *at that point*. For example, if your total wealth was \$100 to begin with, then winning \$100 would represent a doubling of your wealth. Thus, if you won \$100 on the first coin toss, the possible outcomes of the second repetition would be doubled to \$200 and \$0 for heads and tails respectively, matching the doubling of your wealth. If, on the other hand, \$100 only represented 10% of your wealth at the beginning then the amount you would win upon heads would increase only by 10% each time you won.

Clearly, averaging over additive or multiplicative dynamics will give very different results. It is up to us to decide which is more appropriate to the situation. For example, multiplicative dynamics more intuitively applies to situations where results compound, such as compound interest in a market, or population growth, where the number of babies in one generation strongly depends on the number of adults in the previous generation.

It has been shown[10] that, under additive dynamics, the change in wealth that one experiences every time a gamble is repeated is ergodic [11]. This means, under additive dynamics, acting to maximize the ensemble average (ie. the arithmetic average or expected value) of a gamble will also maximize the time average of one's wealth if one repeated the gamble many times. It is worth emphasising that, in the same way that ensemble averaging does not require a real ensemble of people with whom you share your profits in order to be valid, time averaging does not *in practice* require the gamble to be repeated an infinite number of times in order to be useful. Time averaging and ensemble averaging utilize these hypothetical scenarios in order to give an idea of the general properties of a gamble which one can use to make decisions.

Wang. "Economists' Views on the Ergodicity Problem". In: *Nature Physics* 16.12 [2020], pp. 1168–1168)) and instead apply its insights to finding a *normative* decision-making process.

[10]Ole Peters and Murray Gell-Mann. "Evaluating Gambles using Dynamics". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 26.2 (2016).

[11]An important assumption in this proof is that wealth is not bounded from above or below so that one's wealth can become negative without becoming irreversibly bankrupt. This assumption is key and plausibly does not apply to utility functions, which many people argue should be bounded in at least one direction.

2.4 The Original Position and Veil of Ignorance

Harsanyi[12] considered how rational agents would act when deciding on the most desirable distribution of income throughout society. Crucially, he imagined that the agents had to make this decision from what has later been called the 'original position' or 'from behind the veil of ignorance' - a position from which they are impartial, and do not know which place in society they will later occupy. Harsanyi showed that when placed in this situation, rational agents would maximize the expected utility of society (under certain assumptions about what constitutes 'rational' behaviour).

A similar idea is vividly presented by Carlsmith[13], who asks us to consider gambles where the payouts involve saving the lives of 1000 at-risk people. The two available options are labelled A and B. In option A, it is guaranteed that one (randomly selected) person's life will be saved and the other 999 people will die. In option B, there is a 1% chance of saving all 1000 lives and a 99% chance that all 1000 people die. Note that while the certainty of option A has some appeal, option B has a higher expected value of lives saved. Since the people themselves do not know who will be saved in option A, they have to make the decision between A and B from behind the veil of ignorance. For any one of the people at risk, option A offers a 0.1% chance of survival, and option B offers a 1% chance. Thus, behind the veil of ignorance, each person acting to maximize their chance of survival would prefer option B. If each person at risk was given a vote on which option they wanted, maximizing their chance of survival, they would vote for option B. Indeed, in similar situations behind the veil of ignorance, we would always expect people acting to maximize their chance of survival to choose the option which has the largest expected value of lives saved. While this argument does not generalize to all ethical dilemmas, it does provide a powerful reason for maximizing the arithmetic average in certain situations.

3 Geometric Averaging

We introduced the arithmetic averaging decision criterion AADC, which stated that one should choose the option with the highest arithmetic average of the quantity about which you care. We will contrast this with the *Geometric Averaging Decision Criterion* (GADC):

> **The Geometric Averaging Decision Criterion.** When facing a decision between uncertain options, choose the option which has the highest geometric average of the quantity about which you care.

Before providing justifications for the GADC, we will review an important feature of geometric averaging: the arithmetic mean-geometric mean (AM-GM) inequality. The AM-GM inequality states that the arithmetic mean is always greater than or equal to the geometric mean. For a random variable x, this can be expressed as $\mathbb{E}[x] \geq \mathbb{G}[x]$. Also, both geometric and arithmetic averaging hold that the average of a constant value is that value. This means that for a constant x_0: $\mathbb{E}[x_0] = \mathbb{G}[x_0] = x_0$. This fact, combined with the AM-GM inequality has some important consequences. Firstly, if the GADC recommends taking a particular

[12] John C Harsanyi. "Cardinal Utility in Welfare Economics and in the Theory of Risk-taking". In: *Journal of Political Economy* 61.5 (1953), pp. 434–435.

[13] Joe Carlsmith. "On Expected Utility". 2022.

gamble over accepting a particular guaranteed outcome, then (due to the AM-GM inequality), the AADC will also recommend accepting that gamble. However, the converse is not true. There will be some cases where the AADC recommends accepting a gamble over some fixed outcome and the GADC recommends taking the fixed outcome.

We will now provide some justifications for the GADC.

3.1 Ergodicty and Multiplicative Dynamics

In Sections 2.2 and 2.3, we introduced ensemble averaging and time averaging as ways of assessing the general properties of gambles. We argued that arithmetic averaging represents the amount of money one would receive when averaging profits and losses over a large ensemble. Choosing the gamble with the highest arithmetic average also represents the best strategy for maximizing the time average of one's profits when repeating gambles under additive dynamics.

However, under if a gamble is repeated under *multiplicative* dynamics, this is no longer the case. Under multiplicative dynamics, the change in wealth one experiences when a gamble is repeated is not ergodic. This means that acting to maximize the ensemble average of a gamble (ie. the arithmetic average) will not be the correct strategy for maximizing the time average of one's wealth. To see this consider the following simple example, taken from[14].

Imagine you currently have \$100 in your bankroll and are offered the following gamble. A fair coin will be tossed, and if it lands heads you will be given \$50 (representing a 50% increase in wealth), but if it lands tails you will lose \$40 (representing a 40% decrease in wealth). The arithmetic average of your wealth (and thus, the ensemble average) in this gamble is \$105 which is greater than your initial bankroll and would suggest that the gamble should be accepted. However, let us consider what will happen if the gamble is repeated under multiplicative dynamics. Imagine that the gamble is repeated and every time the coin lands heads, your *total* wealth (whatever it may be at that time) is increased by 50% and whenever the coin lands tails your total wealth is decreased by 40%. We will now investigate the behaviour of your wealth as the gamble is repeated many times. Let w_N indicate your wealth after N repetitions of the gamble. Let n_H be the number of times that the coin lands heads, and let n_T be the number of times that the coin lands tails (so that $n_H + n_T = N$). Then we can write:

$$w_N = \$100 \times (1.5)^{n_H}(0.6)^{n_T} \, . \tag{3}$$

We will now try to find an expression for the 'time-averaged growth rate', that is, the average factor by which your wealth is multiplied each time the gamble is repeated, in the limit that the gamble is repeated an infinite number of times. After N repetitions, your wealth has grown by a factor of $\frac{w_N}{\$100}$. Thus, the average factor per repetition by which your wealth has grown is given by the N^{th} root of this value, in the limit that N goes to infinity. Denoting the time-averaged growth rate as $\bar{r}$, we can write:

$$\bar{r} = \lim_{N\to\infty}\left(\frac{w_N}{\$100}\right)^{\frac{1}{N}} = \lim_{N\to\infty}(1.5)^{\frac{n_H}{N}}(0.6)^{\frac{n_T}{N}} = (1.5)^{p_H}(0.6)^{p_T} \, , \tag{4}$$

[14] Peters, "The Ergodicity Problem in Economics".

where we have taken the step of identifying $\lim_{N\to\infty} \frac{n_H}{N}$ with the probability p_H of the coin landing heads on any given trial (and similarly for p_T, the probability of the coin landing tails)[15]. Thus, for a fair coin where $p_H = p_T = \frac{1}{2}$, we find that $\bar{r} = (1.5)^{\frac{1}{2}}(0.6)^{\frac{1}{2}} \approx 0.949 < 1$. Since the time averaged growth rate for this gamble is less than one, when repeated under multiplicative dynamics your wealth will, on average, decrease. As the number of repetitions tends to infinity, your wealth will tend to zero. Thus, time-averaging under multiplicative dynamics suggests that this gamble is not worth taking, even though the expected value is positive.

This may seem counterintuitive but it can be understood quite simply. Suppose you repeated the gamble twice and the coin came up heads once and tails once. This would lead to your wealth being increased by a factor of $1.5 \times 0.6 = 0.9$ - ie. your wealth would decrease. Similarly, your wealth will decrease in any situation where the number of heads equals the number of tails. Since the probabilities of getting heads or tails are equal, in the long run the number of each will approach equality, so you will eventually lose money. This means that repeatedly accepting the gamble will eventually lead one to go bankrupt, even though it has a positive expected value. We will now explore more generally how the geometric average captures this fact when the arithmetic average does not.

Let us consider a quantity x which is gambled under multiplicative dynamics. The starting value of this quantity is x_0. Let m denote the number of possible outcomes of the gamble. For each outcome i, x will be multiplied by a factor r_i (which is greater than 1 for a 'profitable' outcome and less than 1 for an outcome where x decreases). Each outcome will occur with a corresponding probability p_i. After N repetitions of the gamble, the value of x will be:

$$x_N = x_0 \prod_{i=1}^{m} r_i^{n_i}\,, \tag{5}$$

where n_i represents the number of times that outcome i occurred and $\sum_i n_i = N$. Following the same procedure as earlier, the time averaged growth rate will be:

$$\bar{r} = \prod_{i=1}^{m} r_i^{p_i}, \tag{6}$$

where $p_i = \lim_{N\to\infty} \frac{n_i}{N}$ are the probabilities of each outcome. If we are faced with multiple potential gambles and wish to choose the one with the highest time average, we should choose the option with the highest growth rate $\bar{r}$. It is simple to link this growth rate to the geometric average of x. The i-th possible outcome of our gamble leads to a value of x given by $r_i x_0$. The geometric average of these outcomes is therefore given by:

$$\mathbb{G}[x] = \prod_{i=1}^{m} (r_i x_0)^{p_i} = \left(\prod_{i=1}^{m} x_0^{p_i}\right)\left(\prod_{i=1}^{m} r_i^{p_i}\right) = x_0 \bar{r}. \tag{7}$$

Thus, acting to maximize the geometric average of a gamble is equivalent to maximizing the time-averaged growth rate of the gamble under multiplicative dynamics. Again, we emphasize that the geometric average is simply an equation, and we do not require any

[15] This step can be justified through the law of large numbers.

gamble in question to actually be repeated in order to use the geometric average to make decisions. The aim of imagining many iterations repeated one after the other is to give us an idea of the general proclivity of the gamble.

3.2 The Kelly Criterion

The Kelly Criterion is a strategy used in gambling to decide what size of bet to wager, given betting odds along with knowledge of the 'true' probability of each outcome[16]. It has been applied by professional gamblers in blackjack and sports betting as well as by financial traders. Since the Kelly strategy assumes that one is wagering a fraction of one's current wealth at fixed odds, it naturally finds application in multiplicative dynamics, but it equally applies in situations where the odds change between repetitions. What the Kelly strategy amounts to is wagering an amount such that the time-averaged multiplicative growth rate of one's wealth is maximized, not the arithmetic average of each outcome. As discussed in the previous section, this is equivalent to maximizing the geometric average of the outcome.

The Kelly Criterion has been studied in great detail and some powerful results have been proved about its efficacy. In particular, it has been proved that for repeated bets applying the Kelly strategy will, in the long run, outperform any other 'essentially different' strategy, including expected value maximization [17]. Specifically, it can be shown that if one takes the ratio between the bankroll of an agent using the Kelly strategy and the bankroll of an agent using a different strategy, this ratio will tend to infinity as the number of repeated bets tends to infinity. Furthermore, the Kelly strategy will minimize the time taken for one's wealth to reach any given goal, compared to any other strategy. These claims are true, even if the odds change after each gamble, meaning that the Kelly strategy will eventually outperform expected value maximization in any situation where multiple bets are placed one after the other. Proofs of these claims, along with with further discussion of the applications of the Kelly criterion can be found in Kelly's original paper and later works by Thorp, Breiman and Samuelson[18,19,20].

A key feature of the Kelly criterion is that it never advocates wagering one's entire bankroll even if the odds are very favourable (unless the outcome is a certain win). This makes bankruptcy impossible and is in contrast the strategy suggested by the AADC, which would be to wager as much as possible on any positive expectation bet.

As an example, consider the gamble where a fair coin is tossed and, if it lands heads you receive 10 times the amount you wagered plus the original wager, but if it lands tails you lose the amount wagered. Let us denote your wealth with w and your total initial wealth with w_0. If you wager a fraction f of your wealth on this gamble, the arithmetic average of

[16] John L Kelly. "A New Interpretation of Information Rate". In: *The Bell System Technical Journal* 35.4 (1956), pp. 917–926.

[17] The criteria of strategies being 'essentially different' can be made mathematically precise and is included to rule out strategies which are minor variations on the Kelly strategy.

[18] Edward O Thorp. "The Kelly Criterion in Blackjack, Sports betting, and the Stock Market". In: *Handbook of Asset and Liability Management*. Elsevier, 2008, pp. 385–428.

[19] Leo Breiman et al. "Optimal Gambling Systems for Favorable Games". In: *The Kelly Capital Growth Investment Criterion: Theory and Practice*. 1961.

[20] Paul A Samuelson. "The "Fallacy" of Maximizing the Geometric Mean in Long Sequences of Investing or Gambling". In: *Proceedings of the National Academy of sciences* 68.10 (1971), pp. 2493–2496.

your wealth is

$$\mathbb{E}[w] = \frac{1}{2}w_0(1-f) + \frac{1}{2}w_0(1+10f) = w_0 + \frac{9}{2}f\,. \tag{8}$$

Thus, the AADC would advocate maximizing this value by wagering your entire wealth (ie. $f = 1$) on the gamble, even though this would give you a 50% chance of losing everything. On the other hand, the geometric average of your wealth under this gamble would be:

$$\mathbb{G}[w] = w_0(1+10f)^{\frac{1}{2}}(1-f)^{\frac{1}{2}}\,. \tag{9}$$

This expression is maximized when $f = \frac{9}{20}$, which is the fraction of your wealth which the Kelly strategy would advocate wagering. Thus the GADC would advocate wagering $\frac{9}{20}$ of your current wealth on this gamble instead of the entire bank account. While not wagering your full wealth forfeits some gains if you win, it also protects you from bankruptcy and allows you to try the bet again. It is easy to see how if this bet were repeated, an agent using the AADC and repeatedly wagering his total wealth would eventually get unlucky and lose everything. An agent using the GADC and wagering only $\frac{9}{20}$ of his wealth would not run this risk and would slowly build up a large profit.

This provides a powerful rationale for maximizing the geometric average of utililty. To advocate maximizing the expected value of utility as opposed to the geometric average is to advocate a strategy which, when repeated enough times, will almost guarantee that you have less utility. Advocates of expected value maximization would not dispute this result. They would simply point to the fact that, though in *almost all* worlds, an expected-value-maximizing agent will have less utility, in a small number of worlds this agent would have exponentially more utility than a geometric average maximizer. For example, in one possible world, the coin in the previous example would repeatedly land heads giving the expected value maximizer a wealth of $w_0 \times 10^N$ after N repetitions. Though this will only happen with probability $\frac{1}{2^N}$, for the expected value maximizer, this small chance of a very large reward more than compensates for the high likelihood of bankruptcy.

3.3 The Limit where Costs and Payoffs are Small

It is worth here mentioning a useful link between geometric and arithmetic averaging. Notably, in the limit where the gains and losses of the quantity in question are small compared to one's current value, using the arithmetic average to make decisions provides a good approximation to using geometric averaging. Consider a general gamble with starting utility u_0 and payoffs δ_i which can be positive or negative and each occur with corresponding probability p_i. The condition to accept this gamble (over not gambling) according to the GADC is:

$$1 < \prod_{i=1}\left(\frac{u_0+\delta_i}{u_0}\right)^{p_i} = \prod_{i=1}\left(1+\frac{\delta_i}{u_0}\right)^{p_i}\,. \tag{10}$$

If $|\delta_i|$ are small compared to u_0 this expression can be approximated by a first-order Taylor expansion:

$$\prod_{i=1}\left(1+\frac{\delta_i}{u_0}\right)^{p_i} \approx \prod_{i=1}\left(1+p_i\frac{\delta_i}{u_0}\right) \approx 1 + \sum_{i=0} p_i\frac{\delta_i}{u_0}\,, \tag{11}$$

which leads to the condition

$$\sum_{i=1} p_i \delta_i > 0 \,, \tag{12}$$

which is the same as that advocated by the AADC. Thus, in the limit where costs and payoffs are small, decision-making using the the AADC provides a good approximation to using the GADC.

3.4 A Note on Logarithmic Utility

Agents following the Kelly strategy are often described as maximizing the expectation value of a logarthmic utility function. What this means is that an agent who maximizes the geometric average of his wealth behaves in the same way as an agent maximizing the arithmetic average of the logarithm of his wealth. This can be seen by taking the logarithm of the general expression for the geometric average (using ln to indicate the natural logarithm $\log_e$):

$$\ln\big(\mathbb{G}[x]\big) = \ln\left(\prod_i x_i^{p_i}\right) = \sum_i p_i \ln(x_i) = \mathbb{E}[\ln(x)] \,. \tag{13}$$

Since the logarithm function is monotonically increasing, $\ln(\mathbb{G}[x])$ (and hence $\mathbb{E}[\ln(x)]$) encodes the same ordering over preferences as $\mathbb{G}[x]$. Thus, it is argued that an agent acting to maximize the geometric average of a quantity x will behave the same as an agent acting to maximize the *arithmetic* average of the quantity $\ln(x)$. This is normally framed as maximizing the arithmetic average of a *logarthmic utility function.* Recall that a utility function is a function which maps the quantity (x) to a number representing its 'utility', or value. While this is trivially true mathematically, it is worth exploring some subtleties regarding this equivalence. We will provided a few reasons why one should be wary of using geometric averaging and logarithmic utility interchangeably.

Firstly, since the logarithm is only defined for positive nonzero numbers, equation (13) (and therefore the logarithmic utility interpretation) is only valid when none of the outcomes x_i are zero. The geometric average, on the other hand, is well-defined and always equal to zero whenever any of the outcomes are zero. The class of cases where the logarithmic utility interpretation is valid is therefore smaller than the class of cases where the geometric average is valid. This provides one reason for framing decisions in terms of the geometric average instead of the logarithmic expectation.

Secondly, framing geometric averaging in terms of a logarithmic utility function assumes that a logarithmic utility function is the correct one to use. This is a non-trivial assumption as logarithmic utility has several properties representing implicit assumptions which may not fit with other philosophical views. For example, a logarithmic utility function is *concave* meaning that an increase of x by a fixed amount Δx will count for less the larger x is. This represents 'diminishing marginal returns' in x. While this makes intuitive sense when discussing wealth (a billionaire may place less value on a \$100 bill than a bankrupt person), it is less clear that this applies to other, more morally relevant quantities. Is saving one person's life at a time when there are 10 billion other people alive less valuable than saving a person's life at a time when there are 9 billion others? Possibly this is the case, but choosing a concave utility function is an explicit decision that should not be bundled up with the averaging method. Even if one does accept that the utility function should be concave, who is to say that logarithmic utility is the correct way of modelling this, as opposed to some other concave, monotonically increasing function (such as $\frac{1}{1+\frac{1}{x}}$)? The choice of utility function does not come 'for free' when you pick an averaging method.

The problem with representing geometric averaging in terms of logarithmic utility is best exemplified by considering gambles where *utility* itself is at stake. Consider two agents who both agree on the 'true' utility function which should be maximized, but one agent acts to maximize the geometric average and the other acts to maximize the arithmetic average. In situations where there is no uncertainty, both agents will be in perfect agreement about the relative value of different outcomes and will behave identically. But in a gamble where the outcomes have different utilities u_i, each with probabilities p_i, one agent will act to maximize the geometric average of $\prod_i u_i^{p_i}$ and the other will act to maximize the arithmetic average $\sum_i p_i u_i$. It is wrong to say that the agent using geometric averaging has a different utility function since both agents agree on the utilities of every possible state of the world. The agent using geometric averaging has simply used a different criterion to make decisions under uncertainty. Similarly, if an agent decided that utility really was logarithmic in wealth, so that $u = \ln(\text{wealth})$ she would then still be faced with the decision 'should I maximize the geometric or arithmetic average of utility?' If she decides to maximize the geometric average it seems odd in this situation to say that she is *really* just maximizing the arithmetic average of $\ln(\ln(\text{wealth}))$. To claim that maximizing the geometric average is simply maximizing the arithmetic average of the logarithm is to assume somehow that all behaviour must be framed as maximizing the arithmetic average of some quantity. Such a view takes arithmetic averaging as more fundamental than other forms of averaging. This paper is based on rejecting this idea.

Having explored some justifications for both arithmetic and geometric averaging, we will now compare their implications in a variety of situations.

4 A Gamble with Population

In this section, we will introduce a variation on the simple gamble introduced earlier (and adapted from[21]), but instead of using wealth, we will imagine ourselves to be gambling with the future size of a global population. To do this, we will make a few crude assumptions. We assume that increasing or decreasing the population by any amount has no effect on factors other than the population size. For example, increasing the population does not increase resource scarcity or decrease quality of life. Similarly, decreasing the population does not involve killing anyone or violating anyone's freedom. Furthermore, we assume that increasing the population from an initially small pool of individuals does not cause problems (with for example inbreeding). Finally, we assume that all lives are equally 'good' and are all worth living. Thus, other things being equal, according to 'total' consequentialist views, increasing population is always good. Clearly, we are ignoring subtleties which are important in practice, but here we are primarily interested in comparing averaging methods, so will set these issues aside. This thought experiment is to give an idea of how our different averaging

[21]Peters, "The Ergodicity Problem in Economics".

methods might apply to population ethics dilemmas.

Gambling with population size. Consider a planet containing 1 million humans all living happy, fulfilling lives. One day, a wizard appears to the people of the planet and tells them that he is going to put a curse on them but will give them a choice between two different curses. The first curse will lead the planet's population to stay at 1 million people for the rest of time. It will not grow or shrink. This is an unattractive option to many, who wish to increase the number of happy, healthy people living on the planet, so they ask to hear the second option. The wizard tells them "the second option is as follows: every 30 years (corresponding to one generation) I will toss a coin. If it lands heads, your population will become more fertile, and the next generation will be 50% larger than the previous one. However, if it lands tails the fertility of your population will decrease, so that the next generation is 40% smaller than the previous generation." We will represent these options as gambles.

Gamble 1A: Every generation, the population will either increase by 50%, with probability $\frac{1}{2}$ or decrease by 40% with probability $\frac{1}{2}$.

Gamble 1B: Every generation, the population will stay at the same size of 1 million people, with probability 1.

Assuming that the wizard is telling the truth, which curse should they choose?

4.1 Arithmetic Averaging

If the inhabitants of the planet wish to maximize the number of people who exist, they might choose to use the expected value to compare populations subjected to Gamble 1A and Gamble 1B. The expected value of the population after one generation of Gamble 1A is 1.05 million (0.5×1.5 million $+ 0.5 \times 0.6$ million). This suggests a net increase in the population. Gamble 1B, on the other hand only has one possible outcome, so the expected value of the population after option 1B is pursued is 1 million. Comparing these arithmetic averages, the choice clear: Gamble 1A is better than Gamble 1B for producing more people.

But the government are not only concerned about the next immediate generation. Due to the nature of the wizard's curses, their decision will also impact all future generations. Clearly, choosing Gamble 1B would keep the population stable at 1 million forever. On the other hand, if Gamble 1A was chosen, each generation would have a 0.5 chance of increasing the population by 50% over its previous value, and a 0.5 chance of decreasing the population by 40%. Since the arithmetic average suggests an average increase of the population to 1.05 times its previous size, after n generations the expected value of the population size would be $1.05^n \times$ 1 million. Due to the compounding effect, after 100 generations, the expected value of the population would be over 100 times larger than the 1 million which would be achieved by choosing Gamble 1B (since $1.05^{100} \approx 132$). From the point of view of maximizing the population, expected value averaging clearly suggests that Gamble 1A is preferable.

4.2 Geometric Averaging

Using geometric averaging tells a different story. First, let us consider Gamble 1A. Using equation (2), the geometric average of the population (after a single generation) is given by:

$$\mathbb{G}[x] = (1.5 \times 10^6)^{\frac{1}{2}} \times (0.6 \times 10^6)^{\frac{1}{2}} \approx 0.949 \text{ million}\,. \tag{14}$$

On the other hand, the geometric average of Gamble 1B suggests that the population will stay constant, as we would expect:

$$\mathbb{G}[x] = (1 \times 10^6)^1 = 1 \text{ million}. \tag{15}$$

Another, equivalent way of exploring the time-average behaviour of the population is to look at the time-average population growth rate using equation (6). We find that the average growth rate for Gamble 1A is $\bar{r}_A = 1.5^{0.5} \times 0.6^{0.5} \approx 0.949$ and the average growth rate for Gamble 1B is 1. This suggests that the population will *decrease* when repeatedly subjected to Gamble 1A. The time-averaged growth rates that we derived indicated that, if Gamble 1A is repeatedly chosen, the time-average population after n generations will be $0.949^n \times$ 1 million. In other words, the time-averaged population size will *decrease* exponentially. After 100 generations, when the expected value of the population is over 100 million, the time-averaged growth rate predicts that the population is approximately 0.005 million. After 300 generations, the expected value is over 2 trillion, yet the time average population size is less than one. It is reasonable to interpret a population size of less than one person as extinction of that population. How are we to interpret the discrepancy between the high expected value and low (geometric) time-average?

The difference between the two averages is most easily seen by considering a few possible trajectories of the population size over time if Gamble 1A is chosen. There is one possible trajectory where, the wizard's coin lands heads every time, meaning that the population increases every generation, leading to a total population of 1.5^n million. However, very few of the other possible trajectories get close to this level. Indeed, almost all trajectories eventually lead to extinction. When the expected value is considered, the small possibility of a very large population outweighs the fact that the overwhelming majority of possible trajectories will go extinct, if given long enough.

This points to a more general result: over a long period of time acting so as to maximize the time-averaged growth rate minimizes the risk of extinction compared to any other strategy. This is analogous to of the standard theorem (discussed earlier in the context of the Kelly criterion) which states that, over a sufficiently long period of time, maximizing the geometric average of ones financial returns will 'almost certainly' lead to a higher final wealth than any other strategy.

Thus, the two types of averaging highlight two conflicting values which mostly seem uncontroversial: 1) it is good to increase the number of happy lives, in expectation and 2) it is good to not go extinct. As a crude generalisation, we can say that arithmetic averaging emphasizes point 1) and favours increasing the expected value of the population, indifferent to whether it increases the extinction risk. Geometric averaging, on the other hand, recommends foregoing the chance of an exponentially large population in order to secure a stable population. We will explore the issue of extinction in the next section, using a different, more explicit thought experiment.

5 Extinction

In this section, we will consider gambles which explicitly include extinction of the population. Consider the following thought experiment.

> **Gambling with Extinction.** Consider the same planet, with its population of 1 million people. One day, the wizard offers the people of the planet a different gamble:
>
> **Gamble 2A.** The population will double in size with probability p and the population will go extinct with probability $1 - p$.
>
> If they do not accept the gamble, the population will stay at its current value. The wizard asks them: 'how high would I have to make p in order for you to accept the gamble?'.

If one wishes to maximize the expected value of the population, the answer to the wizard's question is simple. The probability p must be greater than $\frac{1}{2}$ in order for the expectation value to be greater than the population would otherwise be.

Using geometric averaging, the answer is more complicated. The geometric average of the population is $(2 \times 10^6)^p \times (0)^{1-p} = 0$, ie. the geometric average of Gamble 2A is always equal to zero, for any nonzero probability of extinction [22].

This makes sense when we view the geometric average in terms of multiplicative dynamics and repeated gambles: once the population goes extinct, it cannot recover. If there is some finite probability of extinction, then it is bound to happen eventually if the gamble is repeated enough times.

In this sense, geometric average fits our intuitions better than expected value reasoning: few of us would accept a 49% chance of extinction for a 51% chance of creating another fully populated earth, yet this is what arithmetic averaging recommends. However, for gambles of this kind, geometric averaging also yields some unintuitive results.

If there is any nonzero probability of extinction, the geometric average of the gamble will be zero. This is due to the fact that, for any such gamble, the geometric average will include a term equal to 0^{p_0} with $p_0 > 0$ which will multiply the rest of the terms. Thus, a gamble with a 1% chance of extinction and a 99% chance of doubling the population and a different gamble with a 99% chance of extinction and a 1% chance of doubling the population both have a geometric-averaged population size of zero and are thus 'equivalent'. This is deeply unintuitive.

Thankfully, different gambles with a geometric average of zero can be compared by considering gambles with finite outcomes of ϵ and taking the ratio of the two geometric averages, in the limit $\epsilon \to 0$. For example, consider a two gambles with an initial population Q.

> **Gamble 3A.** The population will decrease to ϵQ with probability p and increase to aQ (where $a > 1$) with probability $1 - p$.
>
> **Gamble 3B.** The population will decrease to ϵQ with probability q and increase to bQ (where $b > 1$) with probability $1 - q$.

[22] In the case where $p = 1$ and the geometric average is $(2 \times 10^6)^1 \times (0)^0$, we make use the definition $0^0 = 1$.

The ratio of the geometric averages for these two gambles is:

$$\frac{a^{(1-p)}\epsilon^{p}}{b^{(1-q)}\epsilon^{q}}\,. \tag{16}$$

In the limit $\epsilon \to 0$, this ratio either tends to 0 (if $p > q$) or ∞ (if $q > q$), which implies favouring Gamble 3A when $q > p$ and favouring 3B when $p > q$. In other words, geometric averaging favours taking the gamble with the lower risk of extinction, regardless of the other possible outcomes. If the probability of extinction is the same in both cases (ie. $p = q$), then the above ratio reduces to $\frac{a^{1-p}}{b^{1-p}}$ which implies Gamble 3A is preferable when $a > b$ and Gamble 3B is preferable when $b > a$, which matches intuitions. While this does match risk-averse intuitions it can lead to potentially counter-intuitive suggestions as illustrated in the following example.

> **Population Reduction versus Extinction.** You live in a a world containing a population of size Q. You are asked to choose between certainty of a large reduction in the population and the small probability of extinction. These options are represented by the following gambles.
>
> **Gamble 4A.** With probability 1, the the population will be reduced to 0.1% of its original size.
>
> **Gamble 4B.** The population will go extinct with probability $10^{-100000}$ and will increase to however high you would like to go with probability $1 - 10^{-100000}$.

While arithmetic averaging favours Gamble 4B, geometric averaging will always favour Gamble 4A, regardless of how small the probability of extinction is made in Gamble 4B and regardless of how high 4B promises to increase the population if successful. Depending on one's views this is not necessarily a bad thing, but is somewhat counterintuitive, especially for very low chances of extinction. Furthermore, a geometric average maximizing agent will pay any finite cost which does not cause extinction in order to reduce the probability of extinction, even by an arbitrarily small amount. Such preference orderings are a violation of the Von Neumann-Morgenstern Axiom of Contintuity which we repeat here for convenience.

> **Von Neumann-Morgenstern Axiom of Continuity.** For three gambles X, Y, Z where $X \prec Y \prec Z$, there exists some probability $0 < p < 1$ such that $XpZ \sim Y$. In other words, there is some probability p such that one will be indifferent between the guarantee of option Y, and a gamble offering you X with probability p and Z with probability $1 - p$.

The GADC violates this axiom as there is no nonzero probability of extinction which one would accept over a guaranteed loss of population, regardless of how small the probability of extinction is made. Specifically, identifying X with extinction, Y with a (non extinction-causing) population reduction, and Z with a population increase, under the GADC there exists no nonzero p such that $XpZ \sim Y$.

Violations of the other VNM axioms can lead one to accept a series of gambles where you are *guaranteed* to be worse off. In this sense, VNM violations are said to be 'exploitable'.

However, unlike violations of the other VNM axioms, violations of the Continuity Axiom are not guaranteed to be exploitable but instead are only exploitable *with arbitrarily high probability*[23]. In this case, when an agent accepts Gamble 4A over 4B, she is accepting a gamble which is almost certain (but not guaranteed) to leave her worse off. Rejecting the continuity axiom is controversial but has been considered and discussed elsewhere[24,25]. However, in the next section, we will explore a situation where expected value reasoning leads to what could be considered an equally contentious recommendation.

6 Small Probabilities of Enormous Successes

In this section we will investigate situations in which one must pay a small, finite cost for a small probability of a large payoff. For clarity, unlike previous sections, we will assume that what is at stake in the gamble is *utility itself.*

> **Gamble 8A.** Utility is currently at value u_0. With probability p utility will increase to $u_0 + \Delta$, and with probability $1-p$, utility will decrease to $u_0 - \delta$ (assume $\Delta, \delta > 0$).

Under what circumstances should one accept this gamble (over doing nothing and keeping one's utility constant)?

Expected value reasoning will lead us to accept the gamble when the payoff is large enough so that $\Delta > \frac{1-p}{p}\delta$. Note that regardless how small p is, or how large the cost δ, one can always find a value of Δ which satisfies this inequality. If the agent has an unbounded utility function, this value of Δ will correspond to a valid payoff. However, if the agent has a bounded utility function (as, for example, the VNM theorem advocates), then the maximum possible utility is a finite value u_{max}, and the minimum possible utility is 0. Under these circumstances, the largest possible payoff is $\Delta = u_{max} - u_0$ and the largest possible cost is $\delta = u_0$. This means that an agent with a bounded utility function following the AADC can not always be induced to accept Gamble 8A by increasing the payoff, since the payoff cannot be increased arbitrarily, but is instead bounded. Nonetheless, for *any* arbitrarily small probability of success p_ϵ, one can always find a starting utility u_0 and cost $\delta \leq u_0$ for which an expected-value maximizing agent will accept Gamble 8A [26]. Thus, we can always find a situation in which an expected value maximizing agent can always be induced to accept an almost-certain loss for an infinitesimal chance of a large gain.

Let us now compare this behaviour to that of an agent who acts to maximize the geometric average of utility. A geometric average maximizer will accept Gamble 8A provided that the payoff is large enough so that

$$\Delta > \frac{u_0^{\frac{1}{p}}}{(u_0 - \delta)^{\frac{1-p}{p}}} - u_0 \,. \tag{17}$$

[23] Gustafsson, *Money-pump Arguments*.

[24] Larry S Temkin. "Worries about Continuity, Expected Utility Theory, and Practical Reasoning 1". In: *Exploring Practical Philosophy: From Action to Values.* Routledge, 2018, pp. 95–108.

[25] Gustaf Arrhenius and Wlodek Rabinowicz. "Value and Unacceptable Risk". In: *Economics & Philosophy* 21.2 (2005), pp. 177–197.

[26] For example, an expected value maximizing agent will always accept a Gamble 8A when $\delta = \frac{u_0}{2}$ and $\Delta = u_0 - u_{max}$ if they start at a utility satisfying $u_0 < \frac{u_{max}}{1+\frac{1-p}{2p}}$

While this inequality shows that geometric averaging agent can also be induced to accept an arbitrarily small probability of success and an arbitrarily high probability of failure, the payoff required is much higher. Suppose the cost of the gamble is a fraction of u_0, given by $\delta = \frac{u_0}{m}$. The above condition becomes:

$$\Delta > u_0\left(\left(\frac{m}{m-1}\right)^{\frac{1-p}{p}} - 1\right). \tag{18}$$

It can be seen that the payoff required for the gamble to be accepted diverges as $m \to 1$ (reflecting the aversion to ruin discussed earlier) and also diverges exponentially with $\frac{1-p}{p}$. Both of these conditions reflect the more prudent, risk-averse nature of geometric averaging and mean that, for a given arbitrariliy small probability of success, the set of gambles accepted by a geometric average maximizer is much smaller than those accepted by an expected value maximizer. This is because the geometric average maximizer demands a much higher compensation for losses than an expected value maximizer.

Note that this is true when comparing arithmetic and geometric averaging for the same utility function, regardless whether it is bounded or unbounded. Furthermore, even offered unbounded utility, a geometric average-maximizing agent would never accept a gamble which had some finite chance of her receiving zero utility. In this way, while not totally immune, an agent following the GADC is much less susceptible to 'Pascal's mugging'-type scenarios[27] or 'fanaticism'[28] .

7 Democracy and the Veil of Ignorance

In Section 2.4 we introduced the 'veil of ignorance' a thought experiment which seemed to justify arithmetic averaging. We will further explore this theme here and consider some further variations on Carlsmith's thought experiment which we repeat here for convenience.

> **Saving Lives.** The lives of 1000 people who want to live are at risk and they will all die unless you do something. You have two available options:
>
> **Gamble 5A.** With probability 1, a single (randomly selected) person's life will be saved and the rest will die.
>
> **Gamble 5B.** With probability 0.01, all 1000 lives will be saved. With probability 0.99, all 1000 people die.

Expected value reasoning and the AADC recommend that we choose Gamble 5B, even though the probability of success is very small, implying the preference ordering $5B \succ 5A$. One argument in favour of this choice is that, if each person did not know whether they would be the one to be saved, 5B would offer them a larger probability of being saved. Gamble 5A offers each person a 0.1% chance of surviving whereas Gamble 5B offers each person a 1% chance of survival. If each person voted on which gamble to accept, caring only about their own probability of survival, they would vote for 5B.

[27]Nick Bostrom. "Pascal's Mugging". In: *Analysis* 69.3 (2009), pp. 443–445.

[28]Hayden Wilkinson. "In Defense of Fanaticism". In: *Ethics* 132.2 (2022), pp. 445–477.

The geometric average of lives saved in Gamble 5A is 1, but the geometric average of lives saved in Gamble 5B is 0, meaning that using the GADC would imply the preference ordering $5A \succ 5B$. Thus, one could imagine a situation where the decision of which gamble to accept is put to a vote and every person whose life is at risk votes for 5B, yet a decision-maker who uses geometric averaging chooses 5A. This is an important criticism of the GADC in situations of this kind.

This is due to the risk-averse nature of geometric averaging. A result of 0 lives saved is regarded as uniquely bad, since one cannot recover from it under multiplicative dynamics. This makes sense when discussing the extinction of the entire population, but does not carry over to cases such as this, where we are concerned with a relatively small group of people, who we presume do not represent every human in existence. The course of action recommended by geometric averaging changes if we consider an identical scenario, but where the people whose lives are at risk represent a small fraction of the total.

> **Saving Lives part 2**. You live on a planet whose population is 1,001,000. One thousand of them will die unless you do something. You have two available options, detailed below. These options affect only the 1000 people whose lives are at risk. The other 1 million people are unaffected by which choice you make.
>
> **Gamble 6A.** With probability 1, a single (randomly selected) person's life will be saved and the rest will die. Thus, the total population will be reduced to 1,000,001.
>
> **Gamble 6B.** With probability 0.01, all 1000 lives will be saved and the total population will stay at 1,001,000. With probability 0.99, all 1000 people die and the population will be reduced to 1,000,000.

These options are the same as 5A and 5B, except for the addition of 1 million extra people who are unaffected by the outcome of the gamble. The addition of extra people whose lives are not at stake does not change the preference ordering when the expected value is used. Calculating the arithmetic average of the total number of people alive in each situation we find 6A has an expected value of 1,000,001 and 6B has an expected value of 1,000,010. This implies the preference ordering $6B \succ 6A$ which is the same as the earlier gamble which did not include the extra 1 million people. This is due to the fact that the equation for expected value is linear in its arguments. One can add the same constant (in this case, one million extra lives) to each possible outcome and this will result in the expected value shifting by the value of that constant. Since each expected value is shifted by the same amount in each case, this does not affect the ordering over outcomes, meaning that the AADC recommends the same preference ordering, regardless of the extra 1 million people. This is not true for geometric averaging.

Previously, the GADC favoured the guarantee of saving one life over the small chance of saving many. But when a larger population are included, this preference is inverted *even if the extra members of the population are totally unaffected by the outcome of the gambles.* Now, the geometric average of the total population (not just those whose lives are at risk) is 1,000,001 for Gamble 6A and ~1,000,010 for Gamble 6B, implying a preference ordering $6B \succ 6A$. When failing to consider the extra million people, geometric averaging favoured saving one life with certainty, but when the extra million people were considered, it favoured taking the gamble with a 1% chance of saving all 1000 people, agreeing with the expected value-based analysis. This reflects the tendency, noted earlier in Section 3.3, for the GADC and AADC to agree in situations where the costs and payoffs represent a small fraction of a larger whole.

Thus, unlike arithmetic averaging, geometric averaging depends on all members of the class about which you care. The decision of how broad you take this class to be can drastically affect which course of action the GADC recommends. This is because, unlike the arithmetic average, the geometric average is nonlinear. We will explore this issue in more detail in the next Section.

8 Rejecting Background Independence

In this section, we return to examining cases where differences in decision-making hinge on the non-linearity of the geometric average, which we touched upon in Section 7. To demonstrate what is meant by this mathematically consider a gamble with n possible outcomes $\{x_i\}$ with respective probabilities $\{p_i\}$. The expected value of this gamble is $\sum_i^n x_i p_i$. The geometric average is $\Pi_i^n x_i^{p_i}$. Compare this gamble to a gamble with the same probabilities, but the value of each outcome is increased by an amount X. Now, the expected value is $\sum_i^n (X + x_i)p_i = X + \sum_i^n x_i p_i$ and the geometric average is $\Pi_i^n (X + x_i)^{p_i}$. By adding a constant value to each possible outcome, the expected value is simply shifted by that value, leaving preference orderings unchanged (as we saw in Section 7). This is what we mean when using the word 'linear' to describe the expected value. On the other hand, the geometric average does not transform in this simple way when each outcome is increased by a constant value. This can lead to preference orderings recommended by geometric averaging changing depending on the existence of parties not affected by the decision. Consider the following scenario.

> **The Captain's Dilemma.** The two million remaining inhabitants of a planet find themselves having to leave due to the expansion of their local star. The population boards starships and splits into two fleets. One million people (half of the population) heads to the centre of the galaxy, and the remaining million heads to one of the spiral arms of the galaxy, many light years away. Each journey is risky, and there is a significant chance that one or both of the fleets will be partially or totally destroyed en route. Halfway through the journey, communications between the two fleets break down. As a result, the captains of each fleet do not know whether the other is still alive. As the fleet heading towards the spiral arm nears its destination, disaster strikes and half of the ships in the fleet are severely damaged by interstellar debris, to the point where they are unable to continue the journey. This puts the lives of 500,000 people (half of the people in this fleet) at risk. The captain's engineers present him with two options for saving those whose lives are at risk. One option will result in saving the lives of 40,000 of those at risk, while leaving the other 460,000 to certain death. This would mean that 540,000 people would survive from his fleet. Alternatively, the captain could try a risky strategy which would either save all of those at risk, or none of them. The captain's options are represented by two gambles. Note that the numbers given in these gambles refer only to those who are in the captain's fleet and do not include the

members of the other fleet, who are hundreds of light-years away and may not even still be alive.

Gamble 7A. Save 40,000 lives with probability 1. This results in 540,00 people from the fleet surviving in total.

Gamble 7B. With probability 0.1, all 500,000 at-risk people are saved, leading to the full 1 million people from the captain's fleet surviving. With probability 0.9, all 500,000 at-risk people die, leaving only 500,000 survivors in the remaining ships.

If the captain decides to use the expected value to make this decision then he will favour Gamble 7B, (which has an expected value of 550,000 people saved from his own fleet, as opposed to Gamble 7A, which has an expected value of 540,000). Note that, in theory, since the captain is trying to maximize the number of surviving humans, he could also take into account the number of people in the other fleet heading toward the centre of the galaxy, light years away. However, due to the linearity of the expected value, the captain's decision is not affected by whether or not the other fleet is still intact or if it has perished. Regardless of the size of the other fleet, Gamble 7B will still have a higher expected value of 'surviving humans' than Gamble 7A.

Geometric averaging however, *is* affected by the size of the other fleet. If the other fleet has perished then the geometric average of surviving humans is 540,000 for Gamble 7A, and $(1 \times 10^6)^{0.1} \times (5 \times 10^5)^{0.9} \approx 536000$ survivors for Gamble 7B. On the other hand, if the other fleet has survived, the geometric average of surviving humans is 1,540,000 for Gamble 7A and $(2 \times 10^6)^{0.1} \times (1.5 \times 10^6)^{0.9} \approx 1,544,000$ for Gamble 7B. In other words, Gamble 7A has a higher geometric average if the other fleet has perished, and Gamble 7B has a higher geometric average if the other fleet is fully intact. If the captain uses the GADC to make his decision, his choice will depend on the state of the other fleet, even though it is many light years away and can have no causal effect on the captain. If the captain does not know whether the other fleet is intact, he could make the the decision by geometrically averaging over his subjective (Bayesian) probability distribution as to whether the other fleet has survived. Presumably, his assessment of the probability of the other fleet's survival depends on his knowledge of the specific route they were taking and any risks, (such as interstellar debris, etc.) associated with that route. However, now his decision between 7A and 7B depends almost entirely on physical details of the galaxy hundreds of light years away. This seems absurd, analogous to Parfit's 'Egyptology' objection to average utilitarianism[29] (which Parfit himself attributes to McMahan[30]). In this argument, Parfit notes that, under average utilitarianism, the decision whether or not to have a child would depend on whether the child would have a higher welfare than the average welfare throughout history, including the welfare of the Ancient Egyptian. But, as Parfit concludes "research in Egyptology cannot be relevant to our decision whether to have children".

[29] Parfit, *Reasons and Persons*.

[30] Jefferson McMahan. *Problems of Population Theory*. 1981.

As Wilkinson[31], has pointed out, this sensitivity to seemingly irrelevant events results from the rejection of 'background independence': the principle that adding the same given outcome to two lotteries should not affect the preference ordering over those two outcomes. While our example of the captain's dilemma is slightly contrived, it points to an important counterintuitive property of geometric averaging. Unfortunately, the rejection of background independence does not just cause problems in contrived cases such as ours, but it can easily infect all decisions. For example, many ethical frameworks endorse placing at least some value on all forms of sentient life, not just human life. These frameworks would suggest that, if aliens exist in other parts of the universe, even if they are so far away as to never interact with us, we should assign them some ethical value. If we use standard expected value reasoning to make decisions about lives on earth, our decisions will not be affected by the existence (or non-existence) of sentient aliens. However, decision-making based on geometric averaging would potentially depend on whether or not aliens existed, in the same way that the captain's decision depended on whether people in the other fleet were still alive. Echoing Parfit's Egyptology argument, one could claim that 'research in astrobiology millions of light years away cannot be relevant to our decisions on earth'.

In one sense, this is very odd: it seems counterintuitive that a decision made on earth regarding human lives should hinge on the conditions throughout the entire universe, including planets millions of light years away which can have no causal effect on us. Yet this kind of decision-making is not as exotic as it might seem. Consider a conservationist working to preserve a colony of rare parrots in the rainforest. It is not strange if her decision-making changes upon finding out that the parrots are not as rare as she initially thought. If another large colony of parrots was discovered, it is not unreasonable for this knowledge to affect her decision-making, even if the parrots exist hundreds of miles from her. While rejecting background independence can be counterintuitive, it also points to some important considerations. Having explained why some might find the rejection of background independence absurd, we will now muster a brief defense and argue that it is not as absurd as it may seem.

Suppose humanity obtained an oracle which would tell us whether there was intelligent life elsewhere in the universe, even if it was so far away that we would never interact with it. It does not seem absurd to this author that the answer given to us by the oracle would affect our decisions. If we found out that we were alone in the universe, would this not lead us to be more conservative, and prioritize our survival? Conversely, if we found out that there were many similar civilizations to our own throughout the universe, maybe this discovery would empower us to take risks with the chance of a big payoff knowing that if they failed, other civilizations may prosper even if we do not? These are the recommendations that follow from geometric averaging and rejection of background independence. Expected value maximization, on the other hand would recommend policies that would be unaffected by the information from the oracle. While various ethical frameworks may disagree on how to respond to the information from the oracle, it seems counterintuitive to claim that what we learn from the oracle should not affect how we behave in any way. However, background independence (and thus expected value reasoning and the AADC) requires that the information from the oracle should not affect our decision-making *at all*. Similarly, if 100 humans die, the AADC will not distinguish between the case where they are the

[31] Wilkinson, "In Defense of Fanaticism".

last remaining humans on earth and the case where they represent a tiny proportion of a total population containing billions of humans. The GADC on the other hand would take the background information about the total human population into account and rank the former case as significantly worse than the latter. This does not seem unreasonable.

9 Society-wide Decision-Making

So far, we have mostly considered situations where a single decision is made either by an individual, or a group of individuals. In this section, we will briefly explore some situations in which a society of individuals each make their own decisions.

We will consider a society of m individuals who each have a separate utility function. At each timestep, each individual will choose whether or not to take a gamble which affects their utility only. Crucially, the results of the gambles that each member takes will all be independent from each other. We will examine their utilities under three different models of society which we will call IN ('Individualism'), PR ('Partial Redistributionism') and TR ('Total Redistributionism'). In Society IN each person's utility changes only as a result of the gambles that he or she takes and their is no opportunity to share utility between individuals. Under 'total redistributionism', after each round of gambles, the total utility of the entire society is divided equally between each of its inhabitants.

First we note that for a large enough population size m, the society becomes a model for ensemble averaging, meaning that the aggregate utility will be well approximated by M multiplied by the expected value of the gamble. In order for the law of large numbers to apply we require that the smallest probability involved in the gamble p_s is large compared to the reciprocal of the population size, so that $p_s >> \frac{1}{M}$. When this condition is satisfied, the ensemble average interpretation holds and the arithmetic average, multiplied by M provides a good approximation to the aggregate utility of society. In this situation, the expected value of the a gamble provides a good approximation to an individual's utility under TR, even when the population is not infinite. We have previously seen that it is possible for a gamble to have a favourable arithmetic average, but an unfavourable geometric average. Thus, if individuals living in a TR society choose forgo a gamble with a large arithmetic average in favour of one with a more favourable geometric average, they (and the rest of society) will be worse off, since the arithmetic average is the best guide to an individual's utility under TR. If each individual decides to maximize their geometric average, both their individual utilities and the aggregate utility of society will be lower. This provides a strong reason for each individual to choose a gamble with the highest expected value. However, another way of looking at this is that total redistribution with a large enough population simply changes all individual gambles to gambles with a single outcome where one gets the expected value with certainty. The geometric average of this situation is positive, provided that the expected value of each individual gamble is positive.

In the case of the individualist society IN, under repeated gambles, the citizens who choose gambles which maximize their geometric average will, over time, with a probability approaching 1, end up with a higher utility than those who chose to maximize the arithmetic average of their gambles. However, (provided m is large enough for the law of large numbers to apply) the aggregate utility of the society will be maximized when each citizen chooses the option with the highest expected value, since the ensemble average model applies. This can seem paradoxical: suppose each individual was presented with a gamble with a positive expected value, but unfavourable geometric average which would be repeated under multiplicative dynamics. For example, suppose the citizens are offered a gamble where utilities would be multiplied by a factor of 1.5 with probability of 0.5 and multiplied by a factor of 0.6 with a probability of 0.5 (analogous to our Gamble 1A). They have the option to accept the gamble or keep their utilities constant. As we saw before, this gamble has an unfavourable geometric average but a favourable arithmetic average. If every citizen took this gamble repeatedly they would, with probability approaching 1 end up with a lower utility. However, the aggregate utility of the population would increase exponentially due to a small number of citizens whose utility would skyrocket. If each citizen used the GADC, applied to their own utility, they would reject the gamble and keep their utility constant. To a total utilitarian who seeks to maximize aggregate utility, this may be a compelling argument against geometric averaging. However, to the vast majority of citizens would would be better off not accepting the gamble this is a compelling reason in favour of geometric averaging. It is interesting to note that, through a program of partial redistribution of utility, the gamble can be effectively modified so that its arithmetic average is unchanged, its geometric average increases.

Suppose each citizen starts with the same utility u_0 and they are each independently offered the gamble above. If each citizen was a geometric average maximizer, they would reject the gamble, even though it has a positive expected value. Each citizen who wins the gamble will gain $0.5u_0$ and each citizen who loses will lose $0.4u_0$. Now let us consider society PR, which has partial redistribution of utility. That is, there is some mechanism to re-distribute some of the gained utility from those who won the gamble to those who lost. Suppose that half of the gains made by the winners were re-distributed evenly to the losers. Since, in the limit of a large population, the number of winners will equal the number of losers, this would mean that, effectively, the gamble is modified so that the winners gain $0.25u_0$ and the losers only lose $0.15u_0$ each with probability 0.5. This new, effective gamble has the same expected value as the un-redistributed gamble, but now has a positive geometric average of $\sim 1.03u_0$. Thus, even partial redistribution can shift the gamble from having an unfavourable geometric average to a favourable one while keeping the arithmetic average unchanged. Thus, redistribution can incentivize geometric average-maximizing citizens to take riskier gambles than they might otherwise, resulting in a higher aggregate societal utility.

It is important to note that our models used here are crude and do not constitute a proper economic analysis. 'Redistributing' utility does not necessarily map on to redistribution of wealth. It could equally map on to a situation where an individual sacrifices her own utility by working hard to create a product which has great benefit to the rest of society. This is not the place for a discussion on the ethics of redistribution in general. This section is simply intended to point out how redistribution in some form can bridge the disagreement between expected value maximization and geometric averaging.

Broadly, the takeaways from this section are as follows. When the population is large, the arithmetic average (multiplied by the size of the population) provides a good approximation to the aggregate utility of the society. Thus, under a regime of total utility redistribution, an individual faced with a decision should choose maximize the expected value of utility in order to maximize both her own welfare, and the aggregate welfare of society. However, choosing to maximize the arithmetic average in a society with no redistribution can lead to a situation where the choice that leads to the highest aggregate utility leads to the majority of citizens having lower utility than they started with. Thus citizens may be incentivized to maximize the geometric average of their own utility, at the expense of a larger aggregate utility. Finally, a system of partial redistribution does not affect the arithmetic average of a gamble (nor the aggregate expected utility) but can be used to change a gamble with an unfavourable geometric average to one with a favourable geometric average.

10 Conclusion

When applied to ethical dilemmas, both arithmetic averaging and geometric averaging can produce results which match our intuitions in some situations and results that seem unreasonable in others. We will briefly recap some of the major advantages and disadvantages of using each averaging method as a normative decision-making guideline.

We start by listing some disadvantages of the GADC relative to the AADC. Firstly, the GADC rejects the VNM axiom of continuity (Section 5). As a result a geometric average maximizer will accept a certain decrease in utility in place of an infinitesimally small probability of a zero utility result. This is often considered to be irrational behaviour.

Another source of criticism of the geometric averaging approach is its rejection of 'background independence' (discussed in Section 8). The decision advocated by the GADC depends not only on the value of the possible stakes of a gamble but also all members of the class which the agent is trying to maximize. If gambling with utility, the preference ordering over gambles according to the GADC depends not only on the possible changes in utility, but also the starting value of utility. If attempting to maximize the number of living beings, the GADC takes into account not only the lives at stake, but also the lives of people who are unaffected by the gamble. This can lead to situations that some might find absurd, since the decision made can depend on any members of the reference class, even if they are widely separated from the gamble and unaffected by its outcome. This behaviour seems less absurd when one realises it is a result of placing value on changes in utility *as a proportion of the current utility* (as one would do in multiplicative dynamics) rather than in isolation. In this sense, the deaths of 100 humans has more significance if they are the last humans on earth than if they are on a planet with billions of others. When considered in this way, the rejection of background independence does not seem unreasonable.

Due to the arithmetic average's natural interpretation as an ensemble average, the geometric average can give undesirable results in situations when a gamble is repeated many times in 'parallel'. For example, selfish agents acting under a veil of ignorance will prefer to maximize the arithmetic average (as opposed to the geometric average) for gambles regarding their own wellbeing (as discussed in Sections 2.4 and 7). Since the veil of ignorance explicitly assumes that their are many possible positions in society one could occupy, the ensemble average and thus the arithmetic average naturally finds application in these situations. Additionally (as discussed in Section 9), when a society is large enough, the arithmetic average provides a good approximation the aggregate utility of society when each member takes a gamble (or equivalently, the utility of each member of society if utility is equally distributed). It is therefore possible for the GADC to advocate gambles which result in a lower aggregate utility of society.

Now we list some of the advantages of the GADC, relative to the AADC. We have already seen that the GADC can advocate behaviour that may seen excessively risk averse in ‘risk of ruin’ situations involving tiny probabilities of outcomes resulting in zero utility. While the GADC will accept a certain decrease in utility in place of an infinitesimally small probability of a zero utility result, the AADC suffers from a symmetrical problem when posed with gambles with tiny probabilities of large successes (known in informally as ‘Pascal’s Mugging’ scenarios and discussed in Section 6). In particular, the AADC advocates accepting an almost certain decrease in utility in order to access an infinitesimally small probability of a very large increase in utility. The GADC is much more conservative in situations like this. The size of the payoff required to accept such a gamble under the GADC grows rapidly as the probability of success goes to zero and the utility resulting from failure approaches zero. To this author the behaviour of the AADC in ‘Pascal’s Mugging’ scenarios is no more absurd than the behaviour of the GADC in ‘risk of ruin’ scenarios.

Additionally, we have powerful results proved in the context of the Kelly Criterion but applicable widely. An agent acting to maximize the geometric average of utility will (with a probability approaching 1 as the number of gambles they take increases) end up with more utility than an agent acting to maximize the expected value. This is true whenever they are repeatedly taking some kind of gamble, even if the gamble changes from iteration to iteration.

Finally, while the AADC finds natural application in situations involving ‘ensembles’ of agents, the GADC finds natural application as an expression as a multiplicative growth rate when a gamble is repeated in sequence (Section 3.1). Arguably, for most ethical decisions the latter is more appropriate. We saw this in Section 4 where a positive-expectation gamble which looked like it would increase the population size actually lead to extinction when repeated enough times. Unless one lives in a specially designed society (such as TR in Section 9), one does not have a large collective of others with whom to socialize profits and losses. If you go bankrupt buying lottery tickets, the person who won the lottery will not automatically share their winnings with you. This fact becomes particularly important when considering ethical decisions which affect the whole, rather than individuals. If one takes a gamble which affects the whole population of earth, there is not a large ensemble of other earths also taking this gamble with whom we can share resources. There is only us.

This can be highlighted by considering a lottery with a 51% chance of doubling the goodness in the universe now and in the future to come and a 49% chance of destroying all goodness in the universe now and forever. A geometric averaging consequentialist would reject this gamble. An arithmetic averaging consequentialist, on the other hand, would be compelled to accept this lottery, which is deeply unappealing to this author. Furthermore, the lottery could be repeatedly offered and the expected value-maximizing consequentialist would accept each time. Even if she successfully doubled the goodness in the universe a few times, eventually her luck would run out and she would destroy everything good in the universe. This is an outcome which geometric averaging helps us to avoid.