




REPLIES

Egyptians, Aliens, and Okies: Against the Sum of Averages

Christian Tarsney , Michael Geruso and Dean Spears

Population Wellbeing Initiative, University of Texas at Austin, Austin, TX, USA

Corresponding author: Christian Tarsney; Email: christian.tarsney@gmail.com

Abstract

Grill (2023) defends the sum of averages view (SAV), on which the value of a population is found by summing the average welfare of each generation or birth cohort. A major advantage of SAV, according to Grill, is that it escapes the Egyptology objection to average utilitarianism. But, we argue, SAV escapes only the most literal understanding of this objection, since it still allows the value of adding a life to depend on facts about other, intuitively irrelevant lives. Moreover, SAV has a decisive drawback not shared with either average or total utilitarianism: it can evaluate an outcome in which every individual is worse off as better overall, even when exactly the same people exist in both outcomes. These problems, we argue, afflict not only Grill’s view but any view that uses a sum of subpopulation averages, apart from the limiting cases of average and total utilitarianism.

Grill (2023) proposes, as a “contender in contemporary population axiology,” the sum of averages view (SAV), according to which the value of a population is given by

$$\sum_t \bar{u}_t, \quad (\text{SAV})$$

where the t are discrete time periods and \bar{u}_t is the average lifetime welfare of people born in period t . Time periods are discretized by SAV as “periods of substantial length, for example one year” (p. 108), and periods in which no one is born have zero value.¹ Because SAV both sums and takes averages, it inherits some of the well-known implications of both total and average utilitarianism. Like total utilitarianism, SAV would consider it a great loss if a long, prosperous future were replaced with an empty future. But like average utilitarianism, SAV would be indifferent between the actual intertemporal population and an alternative in which each life is replaced with two identical lives, lived concurrently.

¹There are some additional complications to Grill’s view. First, he endorses SAV only for evaluating potential future lives with positive welfare. Second, he suggests taking a sum of *moving* averages, associating each year with the average lifetime welfare of all lives that begin in a 30-year period around it (p. 113). But these complications will have no bearing on our comments.

Although Grill is (to our knowledge) the first to explicitly *defend* SAV as the correct population axiology, it is closely related to social welfare functions that are commonly used for policy evaluation (and, to that extent, implicitly treated as correct) in economics. Macroeconomists and climate economists in particular often examine time-period-specific average utilities or time-period-specific average levels of consumption as a stand-in for welfare (sometimes, but not always, weighted by time discounting).² Scovronick *et al.* (2017) is a prominent recent example in climate macroeconomics that explicitly refers to population ethics and considers a (weighted) sum of per-period averages for calculating social welfare.³ The use of SAV-like views in economics is motivated at least in part by both computational convenience and tradition rather than axiological conviction. Still it is all the more important, in light of these influential applications, to evaluate the plausibility of SAV as a population axiology.⁴

Grill defends SAV as an improvement on standard average utilitarianism that can preserve its attractive features while avoiding its drawbacks. Chief among the drawbacks that Grill aims to avoid is the “Egyptology” objection: the counterintuitive implication that, in Grill’s words, “whether or not it is good that a life starts depends on the quality of lives at other times, including the distant past (such as ancient Egypt)” (p. 104).⁵ Since SAV averages an individual’s lifetime welfare only with her contemporaries (other individuals born at roughly the same time), whether her life contributes positively or negatively to the relevant average does not depend on lives that began in the distant past. So SAV escapes at least a narrow version of the Egyptology objection.

In this comment, however, we argue that SAV does not escape the substance of Egyptology – it is vulnerable to slightly modified but similarly weighty versions of the same objection. We also identify a further, decisive problem for SAV that is not shared with either average or total utilitarianism.

Beginning with Egyptology: broadly speaking, an axiology is vulnerable to Egyptology-like objections insofar as it is *non-separable*. Non-separability means that which of two additions to a population is better (e.g., in the simplest case, adding a single person at a particular welfare level or adding no one) can depend on facts about the pre-existing population. Any view that averages over individuals in a population whose size can vary (such as all the individuals born in a certain period) is non-separable in

²Economists in these literatures are not always explicit about the intended demographic interpretations of their models, for instance whether the utility associated with a given period is the average lifetime utility of a generation/birth cohort or the average momentary utility experienced by anyone alive during a certain timespan. They may also make simplifying demographic assumptions like non-overlapping generations. For both these reasons, social welfare functions in economic models do not always unambiguously represent a fully specific population axiology. In particular, the models we mention here are not usually specified in a way that distinguishes between the “sum of lifetime averages within birth cohorts” view that Grill advocates and the “sum of momentary averages” view that he rejects. But where periods are interpreted as representing or corresponding to generations, as is commonly the case, it seems most natural to interpret the social welfare function that sums average utilities within each period as representing Grill’s view.

³Scovronick *et al.* use a stylized macroeconomic model that does not model births explicitly, but instead simply has period-specific representative-agent consumption.

⁴Especially when we consider the long-term future, SAV can disagree very substantially with total utilitarianism, standard average utilitarianism, and (we speculate) any other plausible axiology. Therefore, if it is not a plausible axiology, it should not be used merely on grounds of convenience, since it is not a good proxy or heuristic for the correct axiology.

⁵This objection to average utilitarianism is originally due to McMahan (1981, p. 115); the now-canonical use of ancient Egyptians to illustrate the objection is from Parfit (1984, p. 420).

this sense. The simplest way to avoid all Egyptology-like objections is to embrace full separability, between all pairs of individuals. Given other plausible assumptions (anonymity, completeness, Pareto, and continuity), this yields an additive view in the “critical-level generalized utilitarian” family that includes total utilitarianism, critical-level utilitarianism, and prioritarianism (Blackorby and Donaldson 1984; Blackorby *et al.* 2005; Thomas 2022).

But not all violations of separability beget equally compelling Egyptology-style objections, and we do not wish to claim that anyone who finds the original Egyptology objection troubling is thereby committed to full separability.⁶ For instance, consider a “relational egalitarian” view that penalizes inequality between individuals who stand in certain morally significant relationships (e.g., engaging in economic transactions with one another or belonging to the same political community). This view is non-separable, and will sometimes imply that facts about unaffected individuals are morally relevant to your choices. For instance, perhaps you ought not create a person with a slightly positive life if they will be much worse off than others in their community. But whether one finds this claim plausible or implausible, it is not implausible in the same way as Egyptology. The welfare of ancient Egyptians seems *obviously irrelevant* to the decision whether to have a baby; the welfare of the people to whom that baby would stand in morally significant relations does not.

Whether SAV successfully avoids the force of the Egyptology objection, then, depends on whether membership in the same birth cohort – i.e., being born at approximately the same time – is enough to make one life intuitively relevant to another. We contend that it is not. Consider, for instance, an ancient Mesoamerican deciding whether to have a child. Are facts about the ancient Egyptians (her unreachable contemporaries across the Atlantic) any more relevant to her procreative decisions than to ours? Would it be reasonable for her to conclude, as SAV implies, that it would be bad to have a child if its welfare, though positive, would be below-average for its birth cohort given the prosperity of those distant Egyptians? (Of course, she would have no way of making such an evaluation in her time – before transatlantic travel or radio communication.) Similarly, consider contemporaries of different species. Does the value of having a child today depend on the size and welfare of, say, present-day fish populations? (Would it be reasonable to delay conception in order to avoid a year when periodical cicadas are hatching, so that your well-off child will be part of a smaller birth cohort and therefore have a larger impact on the average?) Finally, consider individuals in distant galaxies, far enough apart that it is physically impossible even to communicate with them in a human lifetime. Are these distant, unreachable aliens any more relevant to our procreative choices than the ancient Egyptians? It seems to us that all of these implications are as implausible as the original Egyptology counterexample to average utilitarianism.⁷

⁶For a useful discussion of the extent to which anti-Egyptology intuitions support full separability, see §§4.1–4.2 of Thomas (2022).

⁷If the lesson of Egyptology is indeed that (at least) any two lives that stand in no morally significant relationship to one another should be treated separately, then this presents a difficulty not just for SAV but for any view that involves summing over subpopulation averages. This is because the relationships between individuals that are intuitively morally significant are, in general, not transitive. (There may be a chain of morally significant relations connecting you to an ancient Egyptian, without you yourself standing in any morally significant relation to that ancient Egyptian.) They are therefore not equivalence relations, and cannot be used to partition a population into subpopulations.

Table 1. Same-person change that makes every person worse off, but that SAV evaluates as an improvement

Outcome A		Outcome B	
Period 1	Period 2	Period 1	Period 2
Person <i>i</i> born at 11		Person <i>i</i> born at 10	
Person <i>j</i> born at 21		Person <i>j</i> born at 20	
Person <i>k</i> born at 31		Person <i>k</i> born at 30	
$\bar{u}_1 = 26$	$\bar{u}_2 = 11$	$\bar{u}_1 = 30$	$\bar{u}_2 = 15$
SAV = 37		SAV = 45	

The possibility of very large spatial distances between populations highlights a second problem for SAV, which Grill does not consider: special relativity tells us that the time-ordering of events that are spacelike-separated (outside each other’s past/future light cones) is not an objective matter, but depends on a choice of reference frame – observers moving at different speeds will disagree. The same is true of the duration (temporal distance) between timelike-separated events. But SAV is sensitive to these facts about time-order and duration, and its ranking of populations can therefore depend on one’s choice of reference frame. SAV must therefore claim either that axiological facts are velocity-relative, or that there is an axiologically privileged reference frame. Neither of these options is appealing.⁸

Third and finally, SAV has a decisive problem which it does not share with either average or total utilitarianism: it can evaluate a same-people change as an improvement, even though the change makes every person strictly worse off and has no effect on who is born.

Consider the choice between outcome A and outcome B in Table 1. Choosing B delays the birth of person *j* to period 2 and reduces every person’s wellbeing by 1. Despite harming each person, this change has the effect of *increasing* average welfare in both periods. (This is an example of what statisticians call the “Okie paradox,” named for Will Rogers’ joke that “When the Okies left Oklahoma and moved to California, they raised the average intelligence level in both states.”⁹) As a result, SAV implies that outcome B is better than outcome A. This is a verdict rejected by both average and total utilitarianism, and one that seems utterly implausible.¹⁰ It violates even an

⁸Analogous relativity-based objections can be raised to other views that give normative significance to facts about temporal order and duration – for instance, views that incorporate a pure rate of time preference (Cowan 2007, p. 10) and various time-sensitive views in infinite ethics (Cain 1995, p. 404; Wilkinson forthcoming).

⁹We note that “Okie” has a complicated history. The term was once mainly used as a pejorative for migrant agricultural workers from Oklahoma. But it has since acquired much more positive connotations, which we think make it fit to print. It is today used as a badge of pride by native Oklahomans including senator Elizabeth Warren, country singer Vince Gill, and one of the authors of this note (D. S.), as well as the Oklahoma state government (e.g., at <https://oklahoma.gov/branding/verbal-identity.html>). For more on the history of the term, see for instance Keeping (2015) and Paggi (2016).

¹⁰Grill notes that SAV’s evaluations are affected by changing the time period when someone is born, but does not note the point raised by this example: that SAV can evaluate as an improvement a same-person change that makes every person worse off. As a referee pointed out to us, this objection to SAV has been

extremely weak version of the Pareto principle: if exactly the same individuals exist in two outcomes, and one outcome makes each individual strictly worse off, then that outcome is not strictly better.¹¹ More fundamentally, it violates the idea that persons/welfare subjects are the basic loci of value and of moral concern. SAV instead treats *birth cohorts* as basic units of concern. But birth cohorts simply are not fundamental loci of value. It is not plausible to value their “interests” independent of, and at the expense of, the interests of the individuals who compose them.¹²

Grill’s view is just one of a family of possible views that involve summing over subpopulation averages. Indeed, standard average and total utilitarianism are limiting cases of this family: averagism groups all individuals into the same subpopulation (and takes the trivial sum of a single subpopulation average), while totalism treats each individual as their own subpopulation (and sums the trivial welfare averages in these singleton subpopulations). Any view in this family that, like Grill’s, is intermediate between averagism and totalism will inherit many of the weaknesses of both limiting views. Like

previously raised in an unpublished paper by Toby Ord (*ms*), and a similar objection appears in Broome (1992, p. 118).

¹¹Our argument assumes that it is possible for the same person (in our example, person *j*) to be born at different times. Grill makes the same assumption in the example summarized in his Tables 5 and 6 (p. 111).

But suppose you believe that a person’s time of birth is essential to them, so that “person *j*” in our example must be two different people in outcome A vs outcome B. Still, the implication of SAV is hardly less counterintuitive. The case is then reminiscent to Parfit’s classic non-identity case in which a woman can either have a child now, who will have a life that is worth living but fairly difficult, or delay conception at little or no cost to herself, and later have a different child who will have a significantly better life (Parfit 1984, p. 358). This case is interesting because some *prima facie* plausible ethical premises imply the very counterintuitive conclusion that it is *permissible* to have the earlier, worse-off child, even though this decision yields *little or no benefit to anyone else*. But in our case, SAV implies the much more counterintuitive conclusion that it is *better* to have the earlier, worse-off child, *even though this makes everyone else worse off*.

¹²Grill seems to accept the idea that individuals are the basic loci of value when, in discussing the implication of other average utilitarian views that it is better for people with positive but below-average welfare to die sooner, he writes that this is “an inversion of the desired relationship between individual well-being and population value” (p. 110). Elsewhere, however, he seems to suggest that this “desired relationship” does not hold when considering future lives. For instance: “I believe averaging may capture widespread intuitions to the effect that this class of lives [future lives] has value only collectively or generically: it has no value that some extra future life is lived at some unexceptional positive level, but it does have value that future *generations* live good lives or have good life prospects” (p. 107). And similarly: “The appeal of the Average View lies in its ability to capture the well-being of future people in a collective, generic sense – how well future people will be doing collectively, where it does not matter how many they are at any particular time” (p. 116).

We agree with Grill that there are common intuitions in this vicinity. But we do not think these intuitions are well-captured by the idea that “future lives have value only collectively or generically,” or by Grill’s view in particular. For instance, the intuition that “it has no value that some extra future life is lived at some unexceptional positive level” seems better captured by “the Asymmetry” (the claim that we are never obligated to create additional good lives, though we are prohibited from creating additional bad lives) or by critical-range utilitarianism (a separable view on which adding slightly positive lives makes things neither better or worse) than by SAV, which implies that adding an unexceptional future life *does* have positive value if its welfare is even slightly above the average of its birth cohort, has negative value if it is even slightly below, and has “no value” only when it is precisely average.

In any event, even if it could capture our case intuitions, the idea that future birth cohorts are more basic loci of value or bearers of interests than the individuals that make them up seems deeply implausible in itself. The various features that might be thought to confer non-derivative moral importance (sentience, preferences, the capacity for self-perfection, etc.) arguably belong to certain organized collectivities (clubs, parties, corporations), but not to the mere collection of all individuals born in a certain period.

totalism, as Grill acknowledges, it will be vulnerable to the Repugnant Conclusion (at least insofar as the number of subpopulations is potentially unbounded).¹³ Like averaging, it will be vulnerable to Egyptology-style objections and to the Sadistic Conclusion.¹⁴ Moreover, unlike either limiting view, it will sometimes endorse changes that make everyone worse off (at least insofar as it is possible for the same individual to belong to different subpopulations in different outcomes). We therefore believe that all intermediate views in this family will be implausible – there is no viable halfway point, at least on this spectrum, between averaging on the one hand and totalism on the other.

Financial support. Although this paper received no specific funding, the underlying research received support from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (grant K01HD098313 and P2CHD042849) and from the Musk Foundation, awarded to the Population Research Center at The University of Texas at Austin. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funders, who did not review this research.

Competing interests. None.

References

- Arrhenius, G.** 2000. An Impossibility Theorem for Welfarist Axiology. *Economics and Philosophy* 16, pp. 247–66.
- Blackorby, C., Bossert, W. and Donaldson, D. J.** 2005. *Population Issues in Social Choice Theory, Welfare Economics, and Ethics*. Cambridge: Cambridge University Press.
- Blackorby, C. and Donaldson, D.** 1984. Social Criteria for Evaluating Population Change. *Journal of Public Economics* 25(1–2), pp. 13–33.
- Broome, J.** 1992. *Counting the Cost of Global Warming*. Stroud: White Horse Press.
- Cain, J.** 1995. Infinite Utility. *Australasian Journal of Philosophy* 73(3), pp. 401–4.
- Cowen, T.** 2007. Caring about the Distant Future: Why It Matters and What It Means. *The University of Chicago Law Review* 74(1), pp. 5–40.
- Grill, K.** 2023. The Sum of Averages: An Egyptology-Proof Average View. *Utilitas* 35(2), pp. 103–18.
- Keeping, J.** 2015. Once a Slight, Okie Now is Source of Pride for Many. *The Oklahoman*. URL: <https://www.oklahoman.com/story/news/columns/2015/04/07/once-a-slight-okie-now-is-source-of-pride-for-many/60754921007/>. Accessed 21 March 2023.
- McMahan, J.** 1981. Problems of Population Theory. *Ethics* 92, pp. 96–127.
- Ord, T.** ms. On Ranking Alternatives With Different Populations.
- Paggi, R.** 2016. Okie: An Insult or a Badge of Honor? *The Bakersfield Californian*. URL: https://www.bakersfield.com/opinion/okie-an-insult-or-a-badge-of-honor/article_ece50dd3-0b1f-5d9d-88aa-b434e254bf09.html. Accessed 21 March 2023.
- Parfit, D.** 1984. *Reasons and Persons*. Oxford: Oxford University Press.

¹³The Repugnant Conclusion is the implication that, for any population X , there is a better population Y in which everyone has lives just barely worth living (Parfit 1984, §131, pp. 387–90). SAV implies the Repugnant Conclusion since we can make a population of lives barely worth living unboundedly good just by adding more birth cohorts.

¹⁴The Sadistic Conclusion is the implication that it is sometimes better to add a population X in which everyone has negative welfare rather than a population Y in which everyone has positive welfare to the pre-existing population (Arrhenius 2000). For instance, if the preexisting population consists of 10 people with welfare 10, average utilitarianism implies that it is better to add one person with welfare -1 (resulting in an average of $99/11 = 9$) rather than two people with welfare 1 (resulting in an average of $102/12 = 8.5$). Likewise, if a given birth cohort would otherwise consist of 10 people with welfare 10, SAV implies that it is better to add one person with welfare -1 rather than two people with welfare 1.

- Scovronick, N., Budolfson, M. B., Dennig, F., Fleurbaey, M., Siebert, A., Socolow, R. H., Spears, D. and Wagner, F.** 2017. Impact of Population Growth and Population Ethics on Climate Change Mitigation Policy. *Proceedings of the National Academy of Sciences* **114**(46), pp. 12338–43.
- Thomas, T.** 2022. Separability and Population Ethics. In *Oxford Handbook of Population Ethics* edited by Gustaf Arrhenius, Krister Bykvist, Tim Campbell and Elizabeth Finneron-Burns. Oxford: Oxford University Press, 271–95.
- Wilkinson, H.** forthcoming. Aggregation in an Infinite, Relativistic Universe. *Erkenntnis*.