On Moral Aggregation

Frank P. DeVita, December 2025

To Aggregate or Not to Aggregate?

Varieties of consequentialism assent to some version of the proposition that morality turns on outcomes and that actions producing the most good are best.[1,2] In order to determine which outcomes produce the most good, and so which actions are demanded by morality, consequentialists can center *quantities* of well-being, happiness, utility, or other measures of the good and define permissibility and obligation relative to their maximization. The aim of a moral agent's actions, by consequentialist lights, is to produce the most possible good *qua* the maximizing of relevant measures. Many consequentialists also believe in *pure aggregation,* or the summing of quantities of the good.[3] With maximization of the good as a guiding principle and sums or *aggregates* of well-being, consequentialists can analyze conduct in terms of its goodness or badness for individuals or in and across worlds.

---

[1] Sinnot-Armstrong 2023.

[2] Unger 2025.

[3] Tomlin 2017, p. 235.

In this paper, I argue, parallel to Jake Nebel[4], that we should avoid taking sums of well-being for granted, and further that we ought to question the existence of well-being aggregates and thus their moral force. Although aggregates of well-being may be useful in some research frameworks and consequentialist practical reasoning, I maintain that their usefulness doesn't imply their reality and therefore destabilizes their legitimacy as grounds for moral reasoning about action. To argue this position, I first show how aggregative reasoning loses touch with reality, then supplement Nebel's sharp logical critique of sums of well-being with metaphysical and semantic critiques that cast further doubt on moral aggregates and sums of well-being. I then consider some objections defending moral aggregates and, to sum up, briefly sketch an ethical picture that distinguishes between higher-order representations of the good like moral aggregates and first-order well-being, and give moral priority to the latter.

Pointing to Aggregates

What is an aggregate exactly? Definitionally, an aggregate is "a mass or body of parts somewhat loosely associated with one another"[5] or a thing "formed by adding

[4] Nebel 2023.

[5] Merriam-Webster's Dictionary 2025.

together several amounts of things."[6] An arithmetical sum is an aggregate, and sums of things that aren't numbers are also aggregates, e.g. the combined mass of the planets of the solar system. The "total" wealth, health, or happiness of groups or populations are aggregates often used in social science and public health. An aggregate then is a kind of quantity with a relation to, yet a degree of abstraction from, the things it aggregates. That is, aggregates *combine* like properties of disparate things.

Aggregates might also be thought of physically. A froth is an aggregate of bubbles, for instance. Or is it? Upon moving out of the realm of mathematics and into the world looking for aggregates, we encounter difficulties. Is a building an aggregate of bricks? A human being an aggregate of cells? A mind an aggregate of thoughts? These things look *not* to be mere sums of their parts, as aggregation might suggest. In my view, *something is lost* in taking up an aggregative stance. Further, some things don't admit of aggregation. One could aggregate the wealth of the millennials in New York City but could not aggregate the greenness of the trees in Central Park or the minor chords played in all the jazz clubs on Tuesday. Aggregation therefore isn't always possible nor appropriate.

Exploring the logic of aggregation brings its metaphysical oddity into view. Take a simple case. In athletics, aggregate scoring is used to track performance across

[6] Cambridge Dictionary 2025.

matches or rounds, and to break ties. For instance, two soccer teams might have polar opposite match records, e.g. Rome with 5 wins and 2 losses (record 5-2) and Milan with 2 wins and 5 losses (record 2-5). Team A is the clear champion in this hypothetical league of two teams, but soccer aggregates goals, which blurs this fact. Say that Rome won its five matches by scores of 1-0 and lost its two matches by scores of 1-3 and that Milan won its two matches by scores of 1-0 but lost its five matches by scores of 2-3. By match record, Rome is superior. However, by aggregate goals, the score is 14-7 in favor of Milan, since although they lost on the level of matches, they outperformed Rome on the aggregate goal level. If we judge performance in aggregate, the winningest team by the event-level metric (Rome) is a clear loser by an aggregate one. It's difficult, for me at least, to say in what strong sense Rome is truly better than Milan in the example case, except by appealing to some abstract comparison that extracts and decontextualizes events (goals) the situations (matches) in which they actually occurred and pits them against one another.

However, aggregates are useful. In science, aggregates are required machinery to for empirical discovery. Averages, sums, and equations of all kinds are aggregative functions deployed over data sets to construct inferences and reach conclusions about the natural world. So, in a strong sense, aggregates do attach to reality, although they do involve a degree of mathematical abstraction. Take for example the average temperature of Earth. By taking serial temperature measurements over time intervals,

one can determine trends in temperature fluctuations, and such data tells us that the planet is warming—an empirical fact. Still however, it isn't the case, except by coincidence, that any particular place on the planet earth was the average temperature computed by aggregation for the planet. That is, the average temperature, by virtue of being an average—an aggregate—belongs to *the planet* as a whole, but not necessarily to any of the particular locations that make up its parts.

This is not a problem for science or for philosophy. We just have to be aware that assertions containing aggregates are claiming some fact over and above the individuals it aggregates over. That is, for every individual i in an aggregation (Agg) with some property P, a property of the aggregate P_A is a quantity that does not necessarily ($\Box$) belong to, or is not necessarily attributable to, any of the individuals that compose the aggregate:

$$Agg(Pi_1, Pi_2, \ldots Pi_n) = P_A \quad \text{such that} \quad \sim\Box(A = i_1 \vee i_2 \vee i_n)$$

Correlatively, I wonder if the move of *moral aggregation,* i.e. the extracting and comparing of disparate "units of well-being" in abstract as if they co-occurred, *represents* well-being in any proper sense. It certainly does not represent the well-being of any particular individual because it isolates some value and combines it with others from other individuals. At a minimum, some relevant information about individuals is

lost in the process of aggregating. Moreover, we cannot point to anything individual whose well-being is necessarily represented by an aggregate and know if the representation is congruent with their actual well-being by anything other than mathematical chance. By my lights, we ought to be acutely cautious incorporating such abstract representations of well-being into ethics to make claims about permissibility and obligation when lives hang in the balance.

Controversial Quantification

Consideration of well-being is a defining feature of utilitarian consequentialism, which holds that that one ought to choose actions that maximize it.[7,8,9] For utilitarians, moral facts are facts about individual wellbeing.[10] Utilitarians would, I think, also assert the obligation to maximizing at the interpersonal level generalizes to trolley, rescue, or other cases in which one must choose whom or which group to save. For the utilitarian consequentialist, one is obligated to act in accordance with maximizing well-being and therefore to save the many and the valuable. This reasoning entails a sharp bias toward choices that produce or maintain the most *total well-being*. For the consequentialist then,

[7] Bentham 1789.

[8] Mill 1861.

[9] Sidgwick 1907.

[10] Scanlon 1982, p. 108.

aggregates of well-being do some major moral work. But what, *in reality,* do these aggregates grounding the foundations of utilitarian consequentialism represent?

Jake Nebel thinks that although philosophers often talk as if people's well-being can be summed, that it isn't clear this is the case. Is well-being like beauty, specific times, the hardness of materials, or other things that don't come in summable amounts?[11] Nebel looks to measurement theory to sharpen the question and examine the underlying logic of well-being summation and aggregation.[12] Nebel argues that considering well-being as an aggregate measurement leads to some 'highly controversial commitments about the good' that imply we cannot take such sums for granted, as many philosophers tend to.

Nebel says that summing or aggregating well-being entails more than giving it a number, comparing numbers between individuals, and adding those numbers. He thinks well-being is measured interpersonally (comparatively) and distinguishes between *ordinal* and *interval* scales. *Ordinal scales* capture coarse-grained qualitative differences, e.g. the Likert scale.[13] Ordinal scales can capture *that* something is better or worse for a person, but not by how much. *Interval scales* enable finer-grained, numeric measurements, e.g. Celsius and Fahrenheit temperatures, that can capture by how

[11] Nebel 2023, p. 1074.

[12] Ibid, p. 1075.

[13] Likert 1932.

much something is better or worse for someone. Nebel contends that well-being is measurable on interval scales, and further distinguishes between comparisons of differences *qua* gains and losses, and levels themselves *qua* how well off each person is, highlighting that the challenge with these interpersonal-level measurements of wellbeing is determining what the *meaning* of their sums is.[14]

Nebel reasons that well-being is not a number but rather an *extensive quantity* that can be *concatenated* (○) where the former allows for combining of non-numeric quantities like masses and the latter is a function that combines them additively, e.g. ($a \circ b$), and returns a total.[15] For instance, summing the masses of the planets of the solar system by concatenation has the structure $\Sigma m_{\text{planets}} = m_{\text{Earth}} \circ m_{\text{Mars}} \circ m_{\text{Venus}} \circ \ldots m_{\text{Neptune}}$, which gives meaning to "the mass of the planets in the solar system." Sums of well-being, Nebel argues, can be seen similarly. The tricky part, however, is to understand what concatenation does when it operates over well-being, and, crucially, what the resulting quantity represents, and further, *means*.

Nebel offers that we can think of concatenations of well-being as summing life segments, lives, or propositions. Concatenation of *life segments* takes the total well-being W of a life as an aggregate of its temporal segments, e.g. $\Sigma W_{\text{life}} = W_{\text{infancy}} \circ W_{\text{childhood}} \circ W_{\text{adulthood}}$, which allows for comparison and combination of segments between and

[14] Nebel 2023, p. 1076.

[15] Nebel 2023, p. 1079.

across individuals.[16] Segments of life may be at least as good for a person than they are for another (a ⩾ b in Nebel's notation), and concatenation of these would join segments of like well-being to one another, *fusing*[17] them into one.[18] This requires taking up the perspective of the ideal observer, which is fraught with its own problems, including a theoretically damning blindness to the separateness of persons.[19] For Nebel, there are also logical problems with concatenating life segments, but his major concern points out that thinking in this way about well-being makes us focus on the segments of life as disparate rather than how they hang together and constitute the shape of a life.[20]

Concatenating lives is another option. To do this, unions (∪) of populations are considered as sets ($a \cup b$) based on their goodness, how much well-being populations

[16] Nebel 2023, p. 1083-1084.

[17] Nebel traces this *fusion* operation to Rawls 1971, p. 24 where Rawls discusses how the concepts of the impartial spectator and sympathetic identification allow us to lift the ethical principles we apply to individuals to society: "It is this spectator who is conceived as carrying out the required organization of the desires of all persons into one coherent system of desire; it is by this construction that many persons are fused into one…the impartial spectator is the perfectly rational individual who identifies with and experiences the desires of others as if these desires were his own. In this way he ascertains the intensity of these desires and assigns them their appropriate weight…"

[18] Rawls's language mirrors that used in Truthmaker Semantics (Fine 2017, Fine and Jago 2025) to describe the logical operation *fusion* (⊔), which operates over *states* defined as facts, conditions, or events that are mereologically related and can serve as verifiers of a proposition. Truthmaker semantics can also be given for deontic logic in which states are actions and the operators are permission and obligation, suggesting the interesting possibility that this kind of logic and semantics could be quite useful in ethical reasoning.

[19] Rawls 1971 says later in the same paragraph quoted in note 11: "This view of social cooperation is the consequence of extending to society the principle of choice for one man, and then, to make this extension work, conflating all persons into one through the imaginative acts of the impartial sympathetic spectator. *Utilitarianism does not take seriously the distinction between persons*" (italics added).

[20] Dorsey 2015 looks to be a contemporary source of the "shape of life hypothesis." Nebel cites Kamm and Velleman.

contain, or how good each set of lives is for the people living those lives.[21] The problem with this approach, contends Nebel, is that well-being belongs *to a person*, but concatenating over populations combines that well-being into mass that itself belongs to no one. Nebel attempts to surmount this issue by concatenating lives to one another "back-to-back" and taking a sum of well-being in a world to be the well-being of the life sequence composed of all the lives in that world, but worries that these sequences don't exist in any coherent metaphysical sense and nor are compatible with any natural view of personal identity. By Nebel's lights, another problem for the summations so far is that they assume the goodness of a life depends only on its intrinsic properties. The goodness of a life, however, he maintains, could depend on what happens after or before it (e.g. posthumous disgraces or intergenerational projects) and the lives alongside it (e.g. relationships).

Nebel's final attempt to make sums of well-being sensible focuses on propositions as 'bearers of intrinsic value…expressed by *that*-clauses' understood as sets of possible worlds.[22] The general idea is that a set of propositions of the form 'that I am *x*' are intrinsically good or bad for a person and those propositions, taken together as a set can represent the well-being of an individual semantically. The logical idea is to concatenate these sets of propositions with conjunction (∧) to sum well-being. However,

[21] Nebel 2023, p. 1086-1090.

[22] Nebel 2023, p. 1090.

the conjunction two propositions ($p \wedge q$) might be better or worse for a person in ways the logical or linguistic framing misses, e.g. that one had a happy childhood (p) ∧ that one has an unhappy adulthood (q) might not balance symmetrically in the way the formalism ($p \wedge \neg q$) suggests.[23] Then there is the problem of populations, which Nebel claims warrant using *centered propositions* that pair a person (*i, j…*) with a proposition (*p, q…*) that entails some independent facts because the same proposition can be bad for some and good for others. Nebel finds it doubtful that propositions map so cleanly onto people and that we can suspect the relevant set theoretic manipulations drift far into the abstract away from the original task of understanding the well-being of a population.

These speculative attempts at describing sums of well-being lead Nebel to question whether we need them to do ethics, and what reasons we would have to think they exist.[24] In what follows I'll explore both suggestions, and argue for caution about using aggregates in ethical reasoning for metaphysical, semantic, and epistemic reasons.

Goods and Measures

In some sense it seems intuitive that we can measure something like well-being or happiness, and that those measurements are suitable grounds for moral assertions

[23] Nebel 2023, p. 1091-1092.

[24] Nebel 2023, p. 1095.

about what we ought to do. Many scales and measures for different kinds of subjective and objective well-being have been developed that quantify well-being and guide action, including the Gallup Well-Being Index[25], the CDC Health-Related Quality of Life scale, the Subjective Happiness Scale[26], Life Satisfaction Scale[27], Flourishing Index[28], and other novel measures of "affect," "purpose," "optimism," and "thriving."[29]

Morally speaking, the information collected by these scales looks *prima facie* relevant because they purport to quantify total well-being of people and therefore what we ought to do with respect to the populations measured. For instance, the World Happiness report (which uses the Gallup Well-Being Index) shows that Finland is the most happy country while Sierra Leone is the least.[30] We might conclude therefore that Sierra Leone needs more aid or humanitarian attention than Finland, and so that the morally right thing to do at a population-humanitarian level is to aid the former.

The fluidity with which we can reason about sums or aggregates is clear if we reduce the above humanitarian example to a toy one. Consider:

---

[25] Helliwell et al. 2025, University of Oxford: Wellbeing Research Centre.

[26] Lyubomirsky and Lepper, 1999.

[27] Koivumaa-Honkanen 2004.

[28] VanderWeele 2017.

[29] Lee Kum Sheung Center for Health and Happiness 2025.

[30] Afghanistan is in fact the last country on this list; however, its data is largely incomplete except for life expectancy.

> **Cosmic Storm:** In his explorations of deep space, an astronaut finds a pair of planets like ours with people like us inhabiting them. Planet A has a population of 1 billion and Planet B of 2 billion. Both have equally good, enriched lives and tend to live 100 years. The planetary system is endangered by an incoming lethal radioactive storm that will dissolve the planets' atmospheres and end all life on them. The astronaut's ship has one charge of stored nuclear energy he can use to envelop and protect *one* of the planets. Without action, both planets will perish.

The morally right thing to do, I think, is for the astronaut to save the Planet B with 2 billion people. It would also be the morally right thing for a cosmic creator to create Planet B from scratch given the choice. Why? The consequentialist would say that the reason grounding the moral demand is that well-being is maximized by saving planet B. All things being equal, nonconsequentialists could agree with the choice despite disagreeing with the reasoning on various grounds. Now consider:

> **Cosmic Storm 2:** In addition to Planet A and Planet B, an astronaut discovers another world, Planet C, which also contains 2 billion people and is in the path of the lethal radioactive storm. The people on Planet C live mediocre lives compared to those on the other planets. Their culture is devoid of much art or enjoyment, they eat highly processed foods, and tend to live about 50 years.

Which planet ought the astronaut save now, considering the universe of options? For the consequentialist, the answer remains Planet B because that action maximizes well-being given the options set and the context. However, the nonconsequentialist might now, and in my view rightly, respond differently. How? We can represent the universe of options again as such:

	People	**Well-being (*W*, per person)**	**Σ (sum)**
Planet A	1 billion	1 unit	1 billion *W* units
Planet B	2 billion	1 unit	2 billion *W* units
Planet C	2 billion	.5 units	1 billion *W* units

A nonconsequentialist (or at least one of the kind I am trying to be) should question the assertion that saving Planet B is morally obligatory on grounds entailing sums of well-being given this universe and the act of aggregation. How and why are the lives on Planet B *worth more* than the lives on Planet C? The consequentialist answer will appeal to the sum of wellbeing—an aggregate—at stake being greater on Planet B than on Planet C, and use that aggregate of well-being to ground the moral assertion that it is best for the astronaut to save Planet B. Even assume well-being was measured in an accurate and satisfying way as to truly capture the property, why is it that the mere quantity does so much moral work? The sum grounding the assertion, as exposed by

Nebel, is contentious at best. Non-aggregately, there are an equal number of lives at risk between the two worlds the moral agent deliberates between, and that fact is morally relevant. What seems morally reprehensible is the valuing of lives that are going better over lives that aren't going as well.

Ought we not treat all people with equal consideration, or at least consider their lives of equal value? The well-being aggregator has a thorny path through this question. On an aggregative point of view, we lose our concern for individual lives, and act according to a quantity of well-being, which in one sense is theoretically liberating because it allows us to compare across worlds, but in another, that quantity is wholly artificial and misleading because, echoing Nebel, it does not belong to any one person in the group from which it is abstracted. In some strong sense then, the sum of well-being—the aggregate—purportedly represented by the quantity *doesn't exist.* We cannot look at the actual world (or a possible one for that matter) and point to a person whose well-being we are considering when we talk about aggregate well-being. If reasons for actions are grounded by fictive amounts of well-being abstracted away from actual persons' lives, I contend, we are losing our grip on reality when we reason with aggregates. This should leave us weakly motivated to ground ethical reasoning in sums of well-being, or any other aggregates for that matter.

Objections

The moral aggregate theorist ought to have boiled over by now. I have framed well-being aggregates as illusory quantities that fail to represent well-being belonging to people. However, there are entire fields of research across the sociological, psychological, and health sciences that depend on well-being aggregates to characterize people, the world, obligations, and policies. The aggregate theorist could object that the unreality of moral aggregates is an absurd assertion given the empirical sciences that use aggregate metrics and the effectiveness acting on them to make discoveries. These metrics, the aggregate theorist could further argue, allow us to *know* about how well or badly people's lives are at the population level, which is a powerful and important data point in policy design, resource allocation, intervention decisions, and many other humanitarian, developmental, and health contexts. Without well-being aggregation, she might assert, there would be no way to track phenomena that are highly relevant to specifying reasons for action that aim to serve populations, nor an ability to know if interventions designed to improve people's lives are in fact effective.

This objection is pragmatic. It points to a clear and true usefulness for aggregates for important policy-level decisions that can greatly affect the lives of individuals. These aggregates may indeed track variables like well-being that are morally relevant, especially when many lives hang in the balance of a choice. However, these metrics are

still *pseudometrics*—they cobble together many different and disparate variables, and we can question whether those variables are in principle combinable and therefore if the pseudometric measures anything real.[31] The World Happiness Report, for instance ranks the happiness of countries in terms of quantified measures of "social support", "freedom", "generosity", "positive emotions", "negative emotions" and whether people donated, volunteered, or helped strangers.[32] It's not clear that these variables say all there is to say about what constitutes happiness, that summing over them computes happiness, or that they are combinable. Happiness means different things for different people and different theories, and its definition might *not* entail things like volunteering or donation. If they are to be taken seriously, aggregate measures need to specify both criteria for the combining of their constituent variables and a justification for how and why the components and combination are relevant to the phenomenon or quantity in reality they are purporting to track.

Another objection to the denial of moral aggregates might be that we use *collective language* when we refer to the well-being of both individuals and groups in saying things like "he is doing well" or "the people in Sierra Leone are suffering", respectively. Our referring to well-being in this way means that there is aggregate well-being out in the world, and that we refer to it with these kinds of descriptive locutions.

[31] Liao 2025.

[32] https://data.worldhappiness.report/table

After all, there is something meaningful to say about welfare at social and group levels, and we refer to it all the time, especially in the context of government and policymaking. We can query the welfare of certain socioeconomic, geographic, demographic, and other groups by collecting objective and subjective information about them and acting on it accordingly to try and change—and ideally improve—those objective and subjective measures for the group. Without these well-being aggregates, the moral aggregate theorist can contend, we would not be able represent or track the effects of policies, or furthermore, detect progress, success, or failure in attempts at improving people's lives in meaningful ways.

I have semantic and epistemic responses to this objection. The semantic response is that we regularly use language in ways that don't reflect reality. I might refer to a pet as my friend, my friend as a brother, or to a mathematical solution as elegant. Fictional works refer to fictional worlds containing people, places, things, and events that don't exist outside their text. These and other metaphorical or fictive uses of language are pervasive and coherent, but neither implies that the world is anything like the expressions of the language. Pets cannot participate in friendship in the bilateral sense, my friend is not my genetic sibling, and fictional universes are not coextensive with

ours. Reference, I contend, does not necessarily imply existence.[33] The epistemic objection is an extension of the semantic one. Our referring to a thing doesn't imply our *knowledge* of it. We might be mistaken about the existence or nature of the things we refer to with the language we use, and collective, summed, or aggregate well-being might be a thing we talk about, but about the existence or nature of, but about the ontology which we are mistaken.[34]

First-Order Ethics

While higher order representations of the good may be theoretically and practically useful to "compute" what we ought to do, especially at the population level, this paper has attempted to expose the shaky mathematical, logical, metaphysical, and linguistic grounds on which they rest. This negative overdetermination should garner,

[33] Kripke 1973/2013 (reviewed in Azzouni 2016) develops the general idea is that in fictive contexts, we operate under the pretense that we are *pretending* to use the semantic machinery of language to assert, refer, etc. The pretending, in my view, is in virtue of imagining that language refers to fictive worlds and things, that is, people, places, and objects that don't exist. We might do this by accident too, as when I refer (emptily) to the present King of France by this sentence, even if I believe he exists.

[34] Kripke 1980/1970 shows that this happens with natural kind terms as they relate to scientific discovery. Speakers might refer to some material in the mountains as "gold" and later come to find out that it isn't the material with atomic number 79, but rather pyrite and so would have been referring mistakenly. Putnam 1975 develops this further with Twin Earth, showing that speakers can use the same expressions to refer to different things. This certainly applies to words like "good" and "better" which will mean different things to different people.

in my view, some degree of caution toward the use of higher-order representations to define the good, make decisions about action, form moral judgements, and do ethics.

Without moral aggregates, where are we? In my view, we must return to the world and to the people in it to recover a decent ground for moral reasoning about individuals and populations that doesn't distort our picture of the good or detach us from individuals in the ways I have argued well-being aggregates do. Sums of well-being are alluring, especially for population ethics and consequentialist theorizing. However, if we can't articulate a suitable metaphysical ground for them, nor clearly understand how they relate to the individuals we're supposed to be concerned with in the first place when reasoning morally, there isn't, by my lights, good reason to keep them around at risk of distorting our ethical thinking.

Problems with moral aggregates, I think, flag concerns about what an ethical stance ought to accomplish. First, an ethical stance should allow for fluid reasoning about a wide range of cases. If we must shift stances to handle subsets of cases, as happens when taking up a limited aggregation-style position, a stance might need adjusting. There is nothing *wrong* with such stances, but if one seeks stability against inherent instability and of the chancy world, such approaches will fall short in their proclivity to switch frameworks depending on context and conditions at arbitrarily defined thresholds. Second, an ethical stance should *guide action* in ways that most people would find coherent, especially in hard cases. This, I think, is what an ethical

system's chief responsibility is—to tell us what we ought to do given the complexities of real situations. Third, an ethical stance should be realistic. In the heat of a moment, no one has the time or luxury to ponder thresholds, sum well-being, nor consider how someone might reasonably object to their conduct.

Sums of well-being, moral aggregates, and other kinds of pseudometrics in many ways are *proxies* of the good. They are *higher-order* representations, combinations, and concatenations of a thing that belongs to individuals. I cannot share my well-being with you, though I can act to improve both of ours. In this way, well-being, happiness, welfare, quality of life, and other measures like them are distinctly individual in their relation to the value of a life and how well or poorly it's going. They are *first-order* properties of lives that are immanently morally relevant for any moral agent wondering whether or how to act. They are the currency of the moral exchange, and their modulation by actions plays a fundamental role in objective and subjective assessments of conduct and character of a life.

I think we ought to rather focus on the first-order level, that is, on the lives of individuals and ground action in *actual* rather than possible or abstract well-being. At the interpersonal level, this means considering the impact of our conduct on the well-being of others by positing how our choices and actions can *change* their individual well-being from its present state. Actions generating negative impacts are morally contentious and actions generating positive ones are morally good. What will count as

positively and negatively impactful will turn on facts about the individual or group toward whom our actions are directed and their needs, desires, and preferences. At the population level, we may not be able to reason more finely than in terms of whole lives and considering them as, all things being equal, having equal value. Perhaps this isn't a deep moral principle, yet I believe that adherence to it can cut through the noise of the moral world with edges sharp enough to guide minimally decent actions.

References

Azzouni, Jody (2016) 'Reference and Existence by Saul Kripke' *Mind* Vol. 125, No. 498, pp. 593-598.

'Aggregate' (2025) *Cambridge Dictionary.* Cambridge University Press. Accessed December 2025.

'Aggregate' (2025) *Merriam-Webster Dictionary.* Merriam-Webster. Accessed December 2025.

Bentham, Jeremy (1789) *An Introduction to the Principles of Morals and Legislation*. J. H. Burns and H. L. A. Hart, Eds. Oxford University Press, 1996.

Dorsey, Dale (2015) 'The Significance of a Life's Shape' *Ethics* Vol. 125, No. 2, pp. 303-330.

Fine, Kit (2017) 'A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation' *Journal of Philosophical Logic* 46, pp. 625–674.

Fine, Kit and Jago, Mark (2025) *An Introduction to Truthmaker Semantics.* Manuscript.

Helliwell, John F., Layard, Richard, Sachs, Jeffrey D., De Neve, Jan-Emmanuel, Aknin, Lara B., and Wang, Shun, eds. (2025) *World Happiness Report 2025.* University of Oxford: Wellbeing Research Centre.

Koivumaa-Honkanen, Heli, Kaprio, Jaako, Honkanen, Risto, Viinamäki, Heimo, and Koskenvuo, Markku (2004) 'Life satisfaction and depression in a 15-year follow-up of healthy adults' *Social Psychiatry and Psychiatric Epidemiology* Vol. 39, No. 12, pp. 994-999.

Kripke, Saul A. (2013/1973) *Reference and Existence: The John Locke Lectures.* Ocford University Press, New York, NY

— (1980/1970) *Naming and Necessity.* Wiley-Blackwell, Hoboken, NJ.

Lee Kum Sheung Center for Health and Happiness (2025) 'Well-being measurement' https://hsph.harvard.edu/research/health-happiness/well-being-measurement/. Accessed December 2025.

Liao, Matthew S. (2025) 'Saving the Neutrality Intuition from Broome' *Manuscript.*

Likert, Rensis (1932) 'A technique for the measurement of attitudes' *Archives of Psychology* Vol. 22, No. 140, pp. 5-55. R.S. Woodworth, Ed.

Lyubomirsky, Sonja and Lepper, Heidi S. (1999). 'A measure of subjective happiness: preliminary reliability and construct validation.' *Social Indicators Research* Vol. 46, No. 2, pp. 137-155.

Mill, John Stuart (1861) *Utilitarianism*. Roger Crisp, ed. Oxford University Press, 1998.

Nebel, Jacob M. (2023) 'The Sum of Well-Being' *Mind* Vol. 132, No. 528, pp. 1074-1104.

Putnam, Hilary (1975) 'The Meaning of Meaning' In *Language, Mind, and Knowledge,* Minnesota Studies in Philosophy of Science, VII. University of Minnesota Press, Minneapolis, MN.

Scanlon, Thomas M. (1982). 'Contractualism and Utilitarianism' *Utilitarianism and Beyond,* Amartya Sen and Bernard Williams, Eds. Cambridge University Press, pp. 103-128.

Sidgwick, Henry (1907) *The Methods of Ethics*, 7th Edition. Hackett, 1981.

Sinnott-Armstrong, Walter (2025) 'Consequentialism' *The Stanford Encyclopedia of Philosophy* Winter 2023 Edition. Edward N. Zalta and Uri Nodelman, Eds.

Tomlin, Patrick (2017) 'On Limited Aggregation' *Philosophy and Public Affairs* Vol. 45, No. 3, pp. 232-260.

Unger, Peter K. (2025) *Affirming Altruistic Ethics.* Manuscript.

VanderWeele, T.J. (2017). 'On the promotion of human flourishing' *Proceedings of the National Academy of Sciences* Vol. 31, pp. 8148-8156.