How to properly measure moral universalism

Thomas Boyer-Kassem & Sébastien Duchêne

working paper – February 12th, 2026

Abstract. Enke et al. (2022) have introduced the interesting concept of "moral universalism" as "the extent to which people exhibit the same level of altruism and trust toward strangers as toward in-group members". We argue that concerns arise both regarding the theoretical clarity of the concept and the extent to which it genuinely measures what it purports to measure. We distinguish between four reasons for which some economic considerations are not eliminated as expected, due to confounds related to purchasing power disparities, wealth heterogeneity, and relative socio-economic status, and one reason related to the soundness of the distinction between in-group and out-group in the instructions. We conclude by proposing a revised version of the experimental instructions. Even if the past empirical results would have been unchanged by our proposal, we claim that the revised instructions we propose should be used in the future.

In a groundbreaking paper, Enke et al. (2022) have introduced the interesting concept of "moral universalism" as "the extent to which people exhibit the same level of altruism and trust toward strangers as toward in-group members" (p. 3590) (the concept has been caught up by Enke et al. (2023) and Cappelen et al. (2025)). This construct is operationalized through comparisons involving a wide range of social groups, *e.g.*, family, colleagues, religion, or race. Their concept is distinct from that of "altruism" in that "Universalism is not about being more or less altruistic or trusting, but, rather, about a more uniform allocation of a given altruism, or trust 'budget.'" (p. 3590). As a consequence, their experimental design includes questions that ask respondents to split a hypothetical given amount of money, $100, between two people, "a member of an in-group and a randomly selected stranger" (p. 3591). Thus, the experiment tests this concept of moral universalism and not that of altruism, since the level of altruism is fixed to a constant *by construction*.

The authors define and experimentally investigate three kinds of moral universalism: (i) domestic universalism, a trade-off between a domestic in-group and a domestic stranger, for ten different groups: "extended family, friends of family, neighbors, colleagues at work or school, the same organization (*e.g.*, club), the same age, the same ethnic background or race, the same political views, the same hobbies, and the same religious beliefs" (p. 3594) (*e.g.*, between a random person who lives in the US and shares one's religious beliefs—respondents are US-based—and a random person who lives in the US); (ii) foreign universalism, a trade-off between a domestic stranger and a global stranger (*i.e.*, between a random person

who lives in the US and a random person in the world); (iii) global universalism, a trade-off between a global in-group and a global stranger, for five different groups, *i.e.*, with same language, same religious beliefs, same ethnic background, same values, and same occupation (*e.g.*, between a random person in the world who shares the same language and a random person in the world).

Now comes the crucial point on which this Comment focuses. The authors are careful to ensure that the way respondents split the sum of $100 is only affected by the various groups the individuals belong to or do not belong to, and not by their relative monetary wealth or standard of living. Indeed, one may reason that a random person living in the US is on average richer than a random person in the world (see, *e.g.*, https://databank.worldbank.org/source/wealth-accounts/preview/on), so an egalitarian may want to give less to the former and more to the latter because of this wealth inequality, and not because they belong to different groups (which is, again, what the paper aims to investigate). To address this concern, the authors ask the respondents to assume that both recipients are "equally rich" (in the main text, p. 3594 and 3601). The instructions in the experiment read: "Please assume all individuals below have the same income" (Appendix, p. 27). We argue that the instruction to assume both recipients "have the same income" does not effectively eliminate economic considerations as expected, and so for several reasons.

First, a given level of income in the United States implies a vastly different economic reality than the same nominal income in most other countries. Price levels vary considerably across the globe, and what $50 can buy in Ethiopia is not comparable to what it can buy in the U.S. In other words, equal nominal income does not imply equal consumption capacity. As such, asking respondents to assume that both recipients have "the same income," while keeping the transfer amount fixed in U.S. dollars, creates a structural asymmetry in real purchasing power. This undermines the authors' intention to isolate moral universalism from redistributive preferences: participants may attempt to compensate for this imbalance by allocating more to the in-group recipient (US inhabitants)—not out of parochialism, but in an effort to equalize effective material benefit. Thus, a part of the observed data that shows limited universalism (e.g. Fig. 3 p. 3597 in Enke et al. 2022)) could be explained by the fact that respondents have anticipated this discrepancy. Besides, the authors themselves note that "nearly 70% of respondents indicate a strong aversion to inequality" (p. 3601).

Second, the instruction to imagine recipients with the same income may distort how respondents mentally represent the "global" recipient. If income is assumed to be equal across individuals, since very few people in extremely low-income countries (*e.g.*, Uganda) earn as much as an average American, respondents may almost exclude such countries from their mental sample. As a result, they may spontaneously imagine recipients from countries with income levels closer to their own (*e.g.*, France, Canada), which narrows the effective social distance between the "in-group" and the "stranger." This selection effect actually distorts the randomness which is announced in the instructions and dilutes the intended contrast between domestic and foreign recipients. In both cases—whether respondents adjust allocations due to unequal purchasing power or due to perceived proximity in income levels—the interpretation of observed behavior as evidence of limited moral universalism becomes problematic. To address these issues and to properly isolate moral preferences, it would be necessary to control for real purchasing power parity rather than nominal income.

Third, even if we accept the instruction that both recipients have the same income, this does not imply that they possess the same level of wealth. Wealth encompasses not only income flows but also accumulated assets such as property, land, and savings. The correlation between income and wealth varies widely across countries and socio-demographic groups (Pfeffer and Waitkus 2021, Piketty 2014). A household with modest income in a low-income country may have substantial real assets, while an income-equivalent household in a high-income country may be heavily indebted. As a result, assuming income equality does not suffice to eliminate respondents' perception of unequal economic standing. Respondents might rationally infer differences in financial security or long-term stability, leading to allocation patterns based on perceived wealth disparities, rather than group affiliation. Thus, to cleanly measure moral universalism and to focus on the effect of belonging to groups, the experiment should exclude differences in wealth (though see below for discussions on what it could precisely mean), certainly not only in income. So, the phrasing in the main text is correct ("equally rich", p. 3594 and 3601), but that of the instructions in the experiment is not ("Please assume all individuals below have the same income", Appendix, p. 27). Putting all these remarks together, we are led to the conclusion that the instructions in the experiment should ask respondents to imagine two recipients with equal wealth, given purchasing power parity.

Fourth, and importantly, even if two individuals are matched in income and wealth, both adjusted for purchasing power parity, they may still occupy very different relative positions within their respective national contexts. Earning $20,000 a year places you in the lower 10th percentile in the US, but in the very highest ones in Uganda. This positional asymmetry may influence respondents' perceptions of need or deservingness in subtle but significant ways. People often reason about fairness in relative, not just absolute, terms (Festinger 1954). If a foreign recipient is implicitly imagined as affluent within their local context, while the in-group recipient is seen as more modestly situated, allocation decisions may reflect implicit concerns for domestic redistribution rather than preferences over group identity. This form of bias—grounded in comparative social status rather than purchasing power or asset levels—would confound the interpretation of allocation patterns as evidence of limited moral universalism.

In summary, the "same income" instruction fails to neutralize economic considerations. Instead, it allows for persistent confounds related to purchasing power disparities, wealth heterogeneity, and relative socio-economic status, which raises concerns both about the theoretical clarity of the concept of "moral universalism" and its internal validity. It also remains an open question whether improving the formulation of the instructions would leave the empirical results unchanged.

A fifth and final concern relates to the definition of out-groups used to elicit domestic and global universalism. When in-groups constitute a large share of the relevant population, the contrast between an in-group member and a "stranger", drawn from the larger population, tends to disappear. For instance, when respondents allocate money between "a random person who lives in the United States and shares [their] religious beliefs (*e.g.*, a fellow Christian)" and "a random person who lives in the United States," they may recall that about 70% of Americans are Christian (https://news.gallup.com/poll/659339/religious-preferences-largely-stable-2020.aspx), so the 'out-group' person is actually an 'in-group' person with probability 70%! In other words, the task measures the intended concept only with 30% accuracy, and 70%

dilution. It is then no surprise that the measurement of domestic universalism on large groups yields less extreme results than small groups (*e.g.*, family), on figure 3 p. 3597. In such cases, the task may be interpreted as an allocation between two members of nearly the same group rather than a trade-off between an in-group and an out-group member. Then, equal or near-equal allocations would reflect standard within-group fairness norms rather than moral universalism.

Similarly for global universalism about, say, race: for many races (Caucasian, Black, Asian…), a large fraction of the world's population belongs to that race, so there is again a non-negligible probability that a random person in the world is just an in-group. The distinction between a global in-group member and a global stranger is blurred, and the concept of global universalism is again measured with a substantive error – the measured quantity may partly capture within-group egalitarian norms at the global level when the defining group characteristic is sufficiently widespread. By contrast, this concern is much less relevant for foreign universalism, which involves a trade-off between a domestic stranger and a global stranger. In that setting, the contrast between the two recipients is conceptually clear and does not hinge on the size or inclusiveness of a specific in-group, although it could become more salient in contexts where the domestic population itself represents a very large share of the global population (e.g., in the case of China or India). Overall, the concepts of domestic and global universalism are measured with an error rate which increases with the proportion of the in-group in the total population.

Despite our critique of certain aspects of the instructions' wording, we believe that Enke et al. (2022) offer groundbreaking insights into the concept of moral universalism and its experimental measurement. It is therefore worth reflecting on how the instructions could be improved. Drawing a parallel with the instructions presented on page 28 in the Supplementary Materials, we suggest the following revised wording for the case of foreign universalism (with changes highlighted in bold):

"How would you split **the equivalent of** $100 **in real purchasing power** between a randomly-selected person who lives anywhere in the world and a randomly-selected person who lives in the United States? The closer you drag the slider to one individual, the more money you allocate to that individual. Please assume that both individuals **enjoy the same standard of living in their respective countries—including comparable purchasing power and a similar social standing.** Neither individual would find out that it was you who sent them the money."

Our first two concerns are addressed through the use of expressions such as "the equivalent of … in real purchasing power," "same standard of living," or "comparable purchasing power," which help avoid references to fixed amounts and the associated comparison issues. The third concern is resolved by referring to a "same standard of living" rather than to "same income." The fourth is specifically addressed through the mention of "a similar social standing." The revised instructions are formulated in non-technical language and successfully convey the idea that potential economic differences between countries are appropriately taken into account—so that, for a full universalist, a fair split would naturally be 50-50. For the case of domestic universalism, and to address our fifth concern, the phrasing "between a randomly-selected person who lives in the United States and the individual displayed on the right (who is part of a particular social group)"

should be replaced with “between the individual displayed on the right (who is part of a particular social group) and a randomly-selected person who lives in the United States and is not part of that particular social group”. And similarly for foreign universalism.

Our hope is that these changes will be adopted by future users of the experimental design developed by Enke et al. (2022) (and reused in Enke et al. (2023) and Cappelen et al. (2025)) to measure the concept of moral universalism.

References

Cappelen A W, Enke B, Tungodden B (2025) Universalism: global evidence. *American Economic Review* 115(1), 43-76.

Enke B, Rodríguez-Padilla R, Zimmermann F (2022) Moral Universalism: Measurement and Economic Relevance. *Management Sci.* 68(5) 3590–3603.

Enke B, Rodríguez-Padilla R, Zimmermann F (2023) Moral universalism and the structure of ideology. *The Review of Economic Studies* 90(4), 1934-1962.

Festinger L (1954). A theory of social comparison processes. *Human Relations* 7(2) 117–140.

Piketty T (2014) *Capital in the Twenty-First Century.* Cambridge, MA: Harvard University Press.

Pfeffer F T, Waitkus N (2021) The wealth inequality of nations. *American Sociological Rev.* 86(4) 567-602.