Running head: Moral cosmopolitanism

Title: Cosmopolitan morality separates contributing benefits and protecting from harm

Author list: Xuechunzi Bai, Varun Gauri, and Susan T. Fiske

Author institution: Princeton University

Classification: Psychological and Cognitive Sciences

Keywords: Cosmopolitanism, Intergroup morality, Cognitive processes, International policy.

Word count: 7041

Abstract

Global cooperation rests on popular endorsement of cosmopolitan values—putting all humanity on a level with co-nationals, comparative judgments that may tradeoff, even sacrifice, the ingroup’s interests for the interests of all humanity. Moral cosmopolitanism finds support in large, nationally representative surveys from Spain, UK, Germany, China, Japan, US, Colombia, and Guatemala ($N=9,871$). Respondents everywhere distinguish preventing harm to foreign citizens, which almost all support, from redistributing resources, which only about half support. These two dimensions of moral cosmopolitanism, equitable security (preventing harm) and equitable benefits (redistributing resources), predict attitudes toward contested international policies, actual charitable donations, and preferences for mask and vaccine allocations in the COVID response. The dimensions are not correlated with several demographic variables, and only weakly with political ideology. Moral cosmopolitanism also differs from related psychological constructs. Finally, natural language processing probes people’s cognitive associations underlying moral cosmopolitanism (e.g., *world, both*) versus moral patriotism (*USA, first*). Making these global or parochial terms accessible introduces an effective intervention that at least temporarily makes people behave like moral cosmopolitans.

Significance

Contemporary world challenges —pandemic response, climate change, migration, global poverty, and trade policy— require comparative moral judgments between one’s own and other nations. Such choices involve distributing benefits and protecting security for the national ingroup or the whole world. In representative surveys from eight countries, most respondents are willing to sacrifice national well-being to aid foreigners, though more to protect from harms than to redistribute benefits. Cognitive associations (global, inclusive words) reflect cosmopolitan thinking and can elicit cosmopolitan behavior, at least temporarily. Regarding the prevalence and provenance of globally oriented morality, this paper disaggregates two dimensions of moral cosmopolitanism (benefiting versus protecting), distinguishes it from adjacent constructs, demonstrate its relationship to attitudes and behaviors, and proposes a pathway for its enhancement.

/body

Morality beyond borders is hard. In the midst of the pandemic, nations have to decide how much to trade off their own people's interests against foreign nations' interests. Similarly, other contemporary global challenges — climate change, migration, global poverty, and trade rules — require explicitly comparative moral judgments.

The human tendency toward tribal loyalty upsets simple resolutions: Ingroup favoritism complicates moral decision-making at the intergroup level. Although global morality involves many groups, most morality research emphasizes a single group, one's own, construing morality mostly an ingroup problem. Theories focus on morality as an intragroup phenomenon whose primary evolutionary purpose is to reconcile the individual's self-interest with the group's well-being (1-2).

When other groups do join the consideration set, they often take collateral damage from strong ingroup morality (3). For instance, parochial altruism - benefiting ingroup and depriving outgroup - prevails in human behavior, helping the ingroup to survive intergroup competition (2, 4). Likewise, people display parochial cooperation, parochial trust, parochial empathy, and parochial harm sensitivity in intergroup judgments (5-9). Moreover, when situations are direct, zero-sum, and forced-choice, people are much more parochial than cosmopolitan (10).

Morality beyond borders is feasible, nonetheless. Consideration beyond ingroup welfare has precedents. Cosmopolitan philosophers (11-13) argue that one's fellow-citizens' interests should matter no more than those of non-nationals, other things being equal. Corroborating this ideal, evolutionary theory and laboratory research converge on altruism as standing alongside self-interest to drive interpersonal behavior (14-15). At the intergroup level, psychologists propose mechanisms that promote cosmopolitan moral judgments: identification with all

humanity (identity theory; 16) and forces that expand moral boundaries (moral circle theory; 17-18).

Although plausible, none of these cosmopolitanism ideals require ingroup sacrifice to benefit or protect the outgroup. For example, identifying oneself as a world citizen does not exclude being a national citizen; drawing a larger moral circle to include foreign citizens does not exclude fellow citizens from the circle. However, many drivers of conflictual world events demand explicit and direct tradeoffs, such as allocation of international aid budgets or COVID resource distributions, where benefiting or protecting foreign citizens reduces the well-being of one's own citizens, to varying degrees. Given how these tradeoffs complicate many intergroup moral judgments, current understanding of how people make these high-stakes tradeoffs is surprisingly limited. To investigate how people make moral choices between ingroups and outgroups, first, this paper investigates moral cosmopolitanism as overcoming a comparative preference for one's own nation, relative to the world as a whole, in the context of an explicit tradeoff.

Viewing cosmopolitanism as an intergroup choice suggests a second contrast to most approaches. Consistent with the intergroup relations literature, we propose two psychologically distinct aspects: benefit redistribution and harm protection. Benefiting and protecting operate differently. For instance, categorizing people into minimal artificial groups suffices to elicit ingroup favoritism in the allocation of positive outcomes (benefits), but not in allocations of negative ones (harms; 19). As social psychology's intergroup relation theories describe, ingroup love and outgroup hate are not reciprocally related; attachment to one's ingroup does not require hostility toward outgroups (20). Most intergroup discrimination shows preferential treatment of ingroup members rather than direct hostility toward outgroup members (21). Likewise,

intergroup discrimination in cooperation appears motivated by ingroup love rather than outgroup hate (22). In intergroup prisoner's dilemma games, participants most often choose the option that benefits ingroup without affecting the outgroup, rather than choosing the option that simultaneously benefits ingroup and harms outgroup (23).

This is consistent with the normative account of proscriptive and prescriptive morality (3). Proscriptive morality avoids negative outcomes via behavioral inhibition, while prescriptive morality promotes positive outcomes via behavioral activation. In other words, normatively, although people may not help everyone who is in need, they can at least refrain from harming others.

Differentiating the two facets not only improves theoretical clarity but could also be pragmatically instructive. If the harm protection side of moral cosmopolitanism enjoys wider endorsement than benefit redistribution, it may be useful to focus on, and reframe, global justice in those terms -- e.g., protecting all people from the natural disasters associated with climate change, preventing pandemic risk for everyone, lowering the global risk of war and terrorism.

Here, large-scale empirical evidence documents the prevalence and provenance of moral cosmopolitanism, filling gaps from previous arguments often based on American samples. Our findings are based on well-powered nationally representative samples from eight countries in North America, Latin America, Europe, and Asia (total $N = 9,871$). Moral cosmopolitanism correlates with attitudes toward salient world policies and predicts charitable and COVID decisions. For these purposes, moral cosmopolitanism differs from and improves on related psychological constructs (e.g., world identity, moral circle, political ideology). Moreover, identifiable cognitive mechanisms emerge through the analysis of linguistic evidence. From that

comes an intervention that temporarily makes people behave like moral cosmopolitans in benefit redistributions.

Overview of Results

After validating a moral cosmopolitanism scale with two factors, across eight countries, the Equitable Benefit dimension predicts preferences for national or global redistributive choices, with about half of the samples scoring as moral cosmopolitans. In contrast, the second dimension, Equitable Security, approaches ceiling, with most people avoiding differential harm beyond borders. The moral cosmopolitanism scale also predicts attitudes toward controversial world policies (e.g., immigration). Moral cosmopolitanism correlates with but differs from inclusive moral values and political ideology. Identifying psychological processes, Equitable Benefit and Equitable Security highlight relevant social cognitive concepts from linguistic analysis, e.g., reasoning globally. Making the cosmopolitan-related thinking cognitively accessible caused participants to endorse cosmopolitan actions in benefit-redistributive decisions. (Methods follow the results.)

Equitable Benefit and Equitable Security Confirm Distinct Psychological Factors in Moral Cosmopolitanism.

Existing concepts of cosmopolitanism mostly come from sociology (e.g., 24), emphasizing openness to new cultures or enjoyment of travel, but typically not moral judgments. Existing concepts of moral judgments mostly come from psychology (e.g., 25), discussing general moral values but not considering dilemmas that explicitly pit one's group against others. We developed a scale that simultaneously captures cosmopolitanism and moral judgments.

We collected 81 initial items (SI1.1. Item generation) based on existing literature and contemporary concerns, consulted with leading researchers to refine the items, ran pilot studies

with online participants from Western and Eastern cultures ($N = 1,575$). This process reduced the most effective candidate items to 21 (SI1.1. Exploratory and confirmatory factor analysis). We then ran nationally representative samples in eight countries (United States, Guatemala, Columbia, Spain, United Kingdom, Germany, China, Japan; $N=700+$ as the pre-specified sample size per country; $N = 5,772$; SI1.2. Participant demographic information) -- translating and back-translating questionnaires that asked for (dis)agreement on 6-point scales for each item. To extract latent dimensions and test structural (in)variance across national samples, we randomly split responses into exploratory ($N=200$) and confirmatory ($N=500+$) datasets for each country and performed multiple-group factor analyses using maximum likelihood estimation with robust standard errors (SI1.1. Multination factor analysis; 26).

A stable two-dimensional construct showed both country differences and individual differences in all sampled countries (see scale in Table 1; response distributions in Fig.1). The first dimension, *Equitable Benefit* (EB, 5 items), captures the tendency to maximize well-being for broader humanity, even at the expense of one's nation or co-nationals. For example, "Citizens in our country should accept a reduction in their standard of living if doing so increases global equality"; "It is wrong to always prioritize our citizens over people in other countries". Mean responses centered around 3.39, where 6 indicates complete agreement. Empirically, around half (54%) of the participants showed parochial tendencies (below 3.5 on the scale), and 46% showed cosmopolitan tendencies (above 3.5 on the scale) in cross-national benefit redistributions.

The second dimension, *Equitable Security* (ES, 4 items), captures the tendency to protect all people from harm, not just fellow citizens: "It is as bad when terrorist attacks happen elsewhere as when they happen in our nation"; "It is as bad when our military accidentally kills

foreign civilians as when a foreign military accidentally kills our civilians.” Mean responses approach ceiling, with an average of 5.27. Empirically, 95% participants showed cosmopolitan tendencies in cross-national harm allocations (above 3.5 on the scale), and only a very small proportion (5%) of them showed parochial tendencies (SI1.3. Scale descriptive statistics).

Countries showed both similarities and differences. The two dimensions consistently emerge in all sampled countries with satisfactory model fits, indicating the psychological constructs’ generality regardless of cultural or economic contexts (SI1.1. Multination factor analysis). However, we also observe meaningful mean differences: One-way ANOVA revealed significant country differences in Equitable Benefit ($F(7,5764) = 49.894, p < .001$) and Equitable Security ($F(7,5764)=147.08, p < .001$). Post-hoc Tukey HSD tests revealed that the US ($M = 2.93, SD = 1.37$) scored lower on Equitable Benefit than all other countries, and Spain ($M = 4.01, SD = 1.22$) scored higher on Equitable Benefit than all other countries. Speculating, these ranks may reflect the countries’ respective roles in global redistribution (US citizens see their country as donors; Spanish citizens see themselves as a potential recipient of EU aid), or they may reflect differing US-EU attitudes toward the role of public support more generally.

For Equitable Security, Japan ($M = 4.55, SD = .95$) scored lower than all other countries, perhaps because of wariness toward military action arising from the atomic bombings in World-War-II, or generally being more protective of co-nationals rather than other nations. At the other extreme, the Spanish-speaking countries, Guatemala ($M = 5.58, SD = .74$), Colombia ($M = 5.55, SD = .84$), and Spain ($M = 5.66, SD = .68$) scored higher on Equitable Security than the remaining countries.

To summarize, moral cosmopolitanism captures two distinct but related psychological dimensions, Equitable Benefit (resource redistribution) and Equitable Security (harm protection).

The presence of the two dimensions is consistent with our theoretical priors that morality draws on two distinct forms of regulation: approach or promotion (benefit), and prevention or avoidance (respecting security) (27-28). Crucially, the two factors function differently within the intergroup context. Consistent with prior work on asymmetric intergroup evaluations (19,21), people tend to exhibit moral parochialism (ingroup favoritism) in benefit redistribution and moral cosmopolitanism in harm prevention (protect all groups). The fact that the two factors emerge naturally suggests people may think in different ways about benefits versus harms to other groups. The two-factor structure is stable and general across populations, but specific enough to identify differences across countries.

Equitable Benefit Predicts Resource Redistribution; Equitable Security Predicts Avoiding Harm Allocation.

Given that moral cosmopolitanism incorporates two distinct factors of benefiting and not harming outgroups, one would expect that those dimensions predict distinct comparative moral choices. Equitable Benefit should correlate positively with redistributing resources for a global cause relative to a national one, such as charitable donation or COVID19-mask distribution. Equitable Security should be correlated with refusal to harm a non-national to help a co-national, such as sacrificing an immigrant to save citizens in the trolley dilemma (29-30) or sacrificing others to save citizens by monopolizing COVID19-vaccines. Thus, endorsing Equitable Security should correlate with harm protection for all people in the world.

To assess a standard (beneficial) moral behavior, charitable donation, the same online participants from eight countries completed an economic dictator game at the beginning of the study. In addition to their fee, they received \$1, which they could donate to one or divide between two organizations -- Teach for All vs. Teach for America (or their country's equivalent).

The outcome was donations to Teach for All divided by the total amount available to donate. A higher value indicates a morally cosmopolitan preference, while a lower value indicates a morally patriotic (parochial) preference. As a robustness check, we randomly assigned half of the participants to donate on behalf of the research team, while the remaining half had a third option to keep the money for themselves. We ran multilevel regression analyses with individual responses nested within countries allowing for by-country random intercepts and random slopes (SI2.1. Predict charitable donation).

As expected, Equitable Benefit significantly predicted cosmopolitan donations, such that for each unit increase in Equitable Benefit, participants donated 4% ($SE = .01$, $p < .001$) more to Teach for All (Fig 2a). Equitable Security also predicted cosmopolitan donations, but not as strongly. For each unit increase in Equitable Security, participants donated 1% ($SE = .00$, $p = .026$) more to Teach for All. As expected, game-condition main effects emerged: Participants donated 19% ($SE = .03$, $p < .001$) less to Teach for All if they had the option to keep the money (i.e., self-interest). We did not observe interaction effects, meaning self-interest did not moderate the morally patriotic or cosmopolitan donation.

To assess a standard (harmful) moral behavior, participants also completed an international trolley dilemma, making moral judgments on whether it is acceptable to push an immigrant to save five middle-class fellow citizens from a runaway trolley (see Methods and SI2.2. Predict trolley dilemma).

As expected, participants who scored higher on Equitable Security were less likely to push the immigrant ($b = -.37$, $SE = .07$, $p = .001$); Equitable Benefit was unrelated ($b = .024$, $SE = .03$, $p = .433$) (Fig 2b). Given that one central correlate in the trolley dilemma is utilitarianism, we put Equitable Security and instrumental harm, taken from the Oxford Utilitarian Scale (31),

into the same model. We observed a significant interaction between the two ($b = -.12$, $SE = .03$, $p < .001$): For each unit increase in instrumental harm, participants were .49 more likely to condone pushing the immigrant ($b = -.49$, $SE = .04$, $p < .001$). However, if participants were also one unit higher on Equitable Security, they were only .37 more likely to condone pushing the immigrant. In other words, the association between utilitarianism and the international trolley decision attenuated among moral cosmopolitans (SI5.1. Disentangling utilitarianism from cosmopolitanism documents a replication and extension of this experiment).

To explore COVID19-related moral behavior, a preliminary survey addressed commonly discussed choices. Wearing a mask and being vaccinated both provide benefits and protection against harm, but they have different social connotations: “Masks are marketed as tools to help protect others, while the coronavirus vaccine, which may not prevent transmission to others, is perceived as something that helps oneself” (32). Admittedly, when we planned a survey (SI2.3. COVID pre-test), we did not have these predictions. But our findings showed that Equitable Security (endorsing harm prevention for all) was correlated with vaccine distribution. Equitable Benefit (sharing resources that help members of all communities) was correlated with mask distribution. With these priors in mind, we then conducted a twitter-click study (reported below).

First, COVID19 (beneficial) moral behaviors plausibly concern mask redistribution. Since the COVID-19 pandemic began, there have debates surrounding international vs domestically focused policies. In a preliminary test, new online participants in the United States ($N=100$) could sign Petition A (Mask the World) or Petition B (Mask America) via a mocked twitter-click interface (Methods and SI2.4. COVID signing petitions). Respondents indicated endorsement by clicking an interactive retweet button on the screen as many times as they wanted. The petition proposed that the US government and companies should make, sell, and

donate masks to [A: everyone in the world vs. B: our own citizens]. The differences in clicks between the two petitions, which approximate a costly moral behavior (33), was our behavioral outcome. We ran a linear regression with moral cosmopolitanism ratings as the independent variable, the number of click differences as the dependent variable, and political ideology as the control variable.

Masks were supposed to be a resource benefiting other people in the community, especially in the US context. Indeed, Equitable Benefit significantly predicted more support for Mask the World ($b = 22.40$, $SE = 10.31$, $p = .03$), controlling for political ideology ($b = -.52$, $SE = 7.97$, $p = .95$) and Equitable Security ($b = -7.04$, $SE = 13.79$, $p = .61$; neither being predictive). That is, for each unit higher on Equitable Benefit dimension, the participant clicked the Mask the World petition an average of 22 additional times.

COVID19 (harmful) moral behaviors plausibly concern vaccine monopoly. A new set of online American participants ($N=100$) had a choice to sign a different petition: A. Vaccinate the World or B. Vaccinate America. The petition proposed that the US government and companies should make, sell, and donate COVID-19 vaccines to [A: everyone in the world vs. B: our own citizens]. Again, we used the difference in clicks as the behavioral outcome.

We hypothesized that vaccines prevent harm because people are less likely to get infected when vaccinated. As hypothesized, Equitable Security significantly predicted more support for Vaccinate the World ($b = 35.63$, $SE = 13.99$, $p = .01$), controlling for political ideology ($b = 9.97$, $SE = 8.12$, $p = .22$) and Equitable Benefit ($b = 1.65$, $SE = 10.01$, $p = .87$; neither being predictive). That is, for each unit increase in equitable security dimension, participants clicked their support to Vaccinate the World an additional 36 times, on average.

Turning to general policy attitudes: Equitable Benefit correlated with attitudes approving of globally redistributive policies including, for instance, expansive immigration policies (e.g., When jobs are scarce, employers should *not* give priority to people of this country over immigrants, $b = -.37$, $SE = .05$, $p < .001$; and *not* reducing the number of immigrants to our country nowadays, $b = -.21$, $SE = .04$, $p = .001$). Equitable Benefits also predicted positive attitudes toward international organizations (e.g., In general, international institutions do valuable work, $b = .17$, $SE = .03$, $p < .001$) and support for international agreements (e.g., increasing budget support to the UN, $b = .23$, $SE = .03$, $p < .001$; more government budget to assist other countries, $b = .33$, $SE = .04$, $p < .001$).

Equitable Security predicted attitudes toward avoiding active harm including, for instance, criticizing own government and private companies (e.g., boycott companies who make products in other countries just to take advantage of lower safety and health standards, $b = .26$, $SE = .07$, $p = .006$), resisting own country's immoral behaviors, as shown by freely criticizing leaders ($b = .22$, $SE = .07$, $p = .009$) and *not* advocating for the biggest say in deciding world policies ($b = -.14$, $SE = .05$, $p = .025$). It also predicted country-level CO₂ reduction ($b = .35$, $SE = .05$, $p < .001$), but less so personal CO₂ reduction ($b = .09$, $SE = .03$, $p = .005$), perhaps because country-level reduction can reduce active harm to the world, whereas individual-level reduction might be less effective (more items and results in SI2.5. Policy attitudes).

To summarize: The Equitable Benefit factor of moral cosmopolitanism correlated with comparative moral judgments regarding real-world charitable donations, mask redistribution, and redistributive international policies. The Equitable Security factor of moral cosmopolitanism correlated with not sacrificing an immigrant to save fellow citizens, preferences to avoid vaccine monopolies, and opposing policies that actively harm citizens in other countries.

Moral Cosmopolitanism Correlates but Differs from Inclusive Moral Values and Political Ideology.

To understand the mechanisms of moral cosmopolitanism, we examined its potential correlates. Analyses considered plausible demographic associations, as well as other constructs from moral psychology that relate to the greater good, but are not explicitly intergroup tradeoffs or do not distinguish benefits and harms.

In relation to demographics, one might suspect that life experiences, including travel, socio-economic status, age, gender, and political beliefs, were drivers. However, tests generally found slight to null correlations with cosmopolitan donations (SI3.1. Basic demographic correlates), except for political beliefs. Being liberal was slightly and positively correlated with both factors and the actual donations (EB: $b = .27$, $SE = .05$, $p = .001$; ES: $b = .13$, $SE = .02$, $p = .001$; and donation to Teach for All: $b = .03$, $SE = .01$, $p = .016$).

In relation to adjacent psychological constructs, one might also suspect that basic moral values or other psychological constructs explain moral cosmopolitanism. Semi-partial correlation tests found the two factors correlate with distinct sets of psychological constructs (SI1.4. Convergent and discriminant analysis). Specifically, Equitable Benefit does *not* relate to moral foundations of care/harm ($r = .07$ $p = .136$), fairness ($r = .05$ $p = .311$), or purity ($r = -.02$ $p = .738$), but does relate slightly to ingroup ($r = -.16$) and authority ($r = -.21$). In contrast, Equitable Security relates to all fundamental moral motives (harm $r = .29$; fair $r = .28$; ingroup $r = -.15$; authority $r = -.19$; purity $r = -.15$; all $p < .01$).

Equitable Benefit correlates more with beneficial and comparative measures, such as utilitarianism for good, $r = .40$, comparatively more empathy to outgroups than ingroups, $r = .18$ (from 5). Equitable Security correlates more with harmful and generic measures, such as

utilitarianism for bad, $r = -.20$ (from 31), *not* including foreigners in moral concern, $r = -.16$ (from 17), general empathic concern ($r = .36$; from 34), and lower social dominance orientation (dominance dimension, $r = -.34$ and equity dimension, $r = -.31$; from 35). Both factors positively correlated with world identification (EB $r = .25$, ES $r = .25$; from 16; all $p < .001$ otherwise denoted).

In summary, we did not find strong demographic differences in moral cosmopolitanism, with the exceptions of political beliefs, which had small effects. The benefit redistributive factor correlates strongly with ingroup loyalty, consistent with the finding that national boundaries do play a role in benefit allocation. The harm protective factor correlates with general moral concerns, consistent with our finding that national boundaries are less relevant in harm protection. These findings support the notion that people may want to benefit their ingroups, but they do not necessarily want to harm outgroups (19, 21).

Equitable Benefit and Equitable Security Highlight Common and Distinct Social Cognitive Concepts from Linguistic Analysis.

Understanding the cognitive mechanisms underlying moral cosmopolitanism could inform interventions to encourage it. Thus, we probed participants' thoughts by content, analyzing free responses.

To extract most salient words from sentences, in the eight-country study, after making their comparative donations, participants read: "Please tell us, briefly, why you decided as you did." This netted 5772 free responses from people's thoughts about their donations. We developed computational linguistic models that predict which words are most important in predicting their donations (Long-Short Term Memory Recurrent Neural Networks with self-attention; 36-37; SI3.2. Linguistic modeling). We cross-validated the relevant word importance

by regressing the attention weights on participants' donations and moral cosmopolitanism scale scores. In multilevel modeling for each keyword, the independent variable was the attention weight for that keyword for each participant; the dependent variable was that participant's scale scores and actual donations; errors were clustered at the country-level. The regression coefficients thus show which words and concepts matter to moral cosmopolitans, relative to moral patriots (SI3.2, Cross-validation).

Three sets of concepts in the free-response answers were particularly predictive of participants' actual donations. These predictive concepts were: group (*world* vs. *our own*), (im)partiality (*both* vs. *first*), and concrete construal of the cause (*education* vs. *charity*). Although cosmopolitans and patriots both think about the benefited group (not surprising since they were making group-level decisions), the default group differs. Moral cosmopolitans thought more about the world (attention to "*world*" predicts the scale, $b = .68$, $SE = .31$, and predicts donation $b = .77$, $SE = .07$). In contrast, moral patriots thought more about their own [country] (attention to "*our*" $b = -2.68$, $SE = .60$ for scale; $b = -.83$, $SE = .14$ for donation; attention to "*own*" $b = -2.73$, $SE = .97$ for scale; $b = -.73$, $SE = .23$ for donation). Second, concepts of impartiality differed. Moral cosmopolitans paid more attention to equal treatment ("*both*": scale $b = 1.23$, $SE = .30$; donation $b = .50$, $SE = .07$), while moral patriots paid more attention to prioritizing or ranking ("*first*": $b = -1.55$, $SE = .35$; $b = -.55$, $SE = .08$). Third, the levels of other-directed concreteness were different. Moral cosmopolitans focused on improving education for recipients ("*education*": scale $b = 18.46$, $SE = 6.96$; donation $b = 5.17$, $SE = 1.63$), whereas moral patriots focused on the action of the person giving away money ("*charity*": $b = -1.97$, $SE = .60$; $b = -.44$, $SE = .14$; all $p < .05$).

Asking directly for five salient words supplemented the ecologically unconstrained sentence analysis, in a lab-oriented, simpler design with American participants, commenced after the pandemic began. After making their click-for-petition decisions, they read: “Please list briefly some thoughts you had while deciding and while retweeting.” We counted the most frequently mentioned words in the Mask the World, Mask America, Vaccinate the World, and Vaccinate America groups.

The top ten words (with frequencies, distinctive words italicized) in Mask the World are: *world* (13), safety (12), *health* (11), *global* (10), helpful (9), help (9), caring (8), masks (6), *everyone* (6), *love* (5). The top ten words in Mask America are: *America* (24), safety (19), mask (19), *American* (15), help (11), *USA* (10), *protect/protection* (10), *home* (9), *first* (8), helpful (7). The top ten words in Vaccinate the World are: help (23), world (22), *everyone* (18), *global* (18), health (16), *vaccine* (14), *compassion* (11), *equality* (10), safety (9), *hope* (9). The top ten words in Vaccinate America are: help (14), America (13), health (9), care (6), world (6), safety (6), *cost* (5), *first* (5), *give* (5), *citizens* (4). An immediate observation is that, again, we see that group (*world, global, everyone* vs. *America/n, USA, citizens*) and (im)partiality (*everyone, equality* vs. *first, home*) are salient constructs for people to making their retweeting decisions. New, here, is a distinction between other-focused moral cosmopolitan sentiments (*compassion, hope, love*) vs. self-interest (*cost, give*). “Safety” appears in all conditions. But “protection” appears more frequently in Mask America condition (SI3.3. Analysis with single tokens).

In summary, linguistic reports of immediate cognitive accessibility were robust predictors of people’s donations. Concepts such as “world, both, education,” as compared to “own country, first, charity,” are especially predictive if a person is going to donate more to an international cause instead of a national one. We found similar sets of linguistic markers among mask and

vaccine distribution petition supporters. Concepts of “world vs. USA” and “both vs. first” appeared frequently. “Protection” was salient for mask distribution.

Making the Cosmopolitan-related Thinking Cognitively Accessible Caused Participants to Endorse Cosmopolitan Actions in Benefit Redistributive Decisions.

Having identified an apparent cognitive mechanism, we wanted to use it to experimentally encourage people to engage in cosmopolitan actions. This experiment examined the causal effects of the three classes of cosmopolitan concepts, identified above, using semantic priming. Following the eight-country studies, we hypothesized that, if concepts like “world, both, education” are the active concepts of cosmopolitan-related thinking, making them cognitively accessible would lead people to engage in cosmopolitan action. Conversely, we thought that concepts like “our own, first, charity” would activate moral patriotism and parochial generosity. We did not intervene in harm protective decisions because they are already endorsed almost at ceiling; instead we focused on the hard decisions – providing benefits to the world at ingroup expense.

For a semantic intervention, New American online participants ($N = 1162$) were randomly assigned to read either moral cosmopolitan or moral patriotic (parochial) concepts. We selected sentences from previous American participants that included at least two of the key concepts. Note that these sentences simply supplied individual reasoning, not behavioral norms or overt persuasion. Crucially, to make sure participants actively engaged in the concepts and thought about them, rather than receiving them passively, we deployed methods from the Elaboration Likelihood Model (38) and Narrative Paradigm (39) to assist the semantic priming task. Participants typed the pros and cons or told a personal story regarding each treatment sentence.

The outcome measure was real-world charity donation (share of donations, out of \$1, to international organizations rather than national ones). To reduce demand effects, we listed six nongovernmental organizations (NGOs: national and international efforts respectively in education, poverty, emergency medicine) (SI4.1, Intervention participants and procedures).

The baseline cosmopolitan donation rate from previous studies was 35.7%, meaning that on average participants donated more to national NGOs than to international NGOs. Our experiment revealed a significant average accessibility effect (Fig 4). Participants treated with moral cosmopolitanism concepts donated 50% to international NGOs, 11.8% ($SE = .02$, $p < .001$); those treated with moral patriotic concepts donated 38%. We did not find engagement-style differences: Participants who wrote pros and cons donated similarly to those who wrote personal stories. Although we tentatively hypothesized that extreme believers (in either direction) might be less influenced by our manipulation than moderate believers, we did not observe such treatment heterogeneity. To rule out the potential confounding effect of support for education NGOs, we tested out-of-domain outcomes by excluding Teach for All/America and found a similar effect ($b = .10$, $SE = .02$, $p < .001$). This test also mitigates concerns related to demand effects (SI4.2. Intervention robustness check).

In summary, confirming the linguistic evidence of correlated accessibility, this experiment demonstrated the causal impact of moral cosmopolitan concepts on international donations. Manipulating participants to think like moral cosmopolitans changed their donations in a resource-allocation context, moving them to help the world even at the expense of their own nation and fellow citizens. Resource allocations are influenced by tribal morality (2,4), but the current experiment shows that interventions could plausibly break down tribal (parochial) boundaries, at least temporarily.

Discussion

This paper investigated what people do when facing direct, comparative moral judgments between their own and other nations. This context mimics the hard tradeoffs that world leaders and citizens increasingly face. For example, should people prioritize giving masks and vaccines to their own citizens, or make the resources equally available to people around the world? Under the broad umbrella of moral cosmopolitanism, defined as overcoming a comparative preference for one's own nation, relative to the world as a whole, in the context of an explicit tradeoff, we found two factors: benefit redistribution and harm protection.

Equitable Benefit, on the one hand, relates to redistributing resources to non-nationals. Consistent with ingroup biases in the moral judgment literature, people in such comparative choices tend to remain parochial. Equitable Security, on the other hand, entails not harming non-nationals. Most people seem to overcome comparative choices on that dimension, endorsing safety for all. The distinction between the two factors is consistent with intergroup relations literature in social psychology, which shows asymmetric valence judgments (19). Benefiting and harming outgroups are independent. On the basis of sizable samples from eight countries in North America, South America, Europe, Asia, we provide evidence for this distinction.

This distinction has theoretical implications. Past research shows mixed results. Some studies find people are mostly parochial, others find people do possess a global morality. Joining the conversation, our results suggest focusing on whether a judgment concerns benefit redistribution or harm protection. They are not symmetrical, especially in intergroup judgments. When people are ready to help the ingroup, we should observe tribal or parochial morality; when people are unwilling to harm outgroups, we should observe cosmopolitan morality. Future

research might draw on this differentiation when assessing intergroup moral judgments and global thinking, more generally.

This differentiation can also be pragmatic because the two dimensions predict distinct sets of policy attitudes and real-world moral behaviors. Equitable Benefit predicts redistributive policies such as immigrant aid, charitable donations, and mask distributions. Equitable Security predicts harm protective policies such as human rights violation, trolley dilemma, and vaccine monopoly. This differentiation can inform expectations about how people will react toward policies that either imply benefitting an in-group or harming an outgroup. When the policy is (or framed as) ingroup beneficial, the Equitable Benefit dimension is likely to be operative in people's minds. When the policy is (or framed as) outgroup harmful, we should expect Equitable Security psychology to be at work.

Throughout these analyses, political ideology only slightly correlates with cosmopolitan judgments; its predictive power diminishes after controlling for moral cosmopolitanism. This indicates that many controversial world policies may transcend partisanship, and may entail deeper psychological underpinnings.

The mechanism underlying moral cosmopolitanism seems to suggest distinct sets of cognitive concepts when people are making comparative decisions. Our finding is consistent with philosophical accounts that emphasize the role of reason in generating global concern (40, 12). Linguistic analysis and experiments found that moral behavior can change by manipulating the framing of global thinking. In the experiment, we did not invoke a larger moral circle, which could both create a common global identity and encourage utilitarian considerations. Linguistic cues given to moral patriots sufficed to increase cosmopolitan donations regardless of participants' original political ideology, age, gender, education level, or place on the social

ladder. This result speaks to prior research that found altruistic behavior unresponsive to framing (41). To be sure, our manipulation changed only temporary behaviors; we do not have data on long-term effects. Probably, people require prolonged exposure to cosmopolitan concepts to overcome their everyday exposure to parochialism. Future research could explore longitudinal experiments.

We focused on the nation as the group boundary of interest because it is psychologically understudied, yet relevant to important political, economic, and social conflicts. Furthermore, moral patriotism is philosophically defensible, at least in comparison to other morally arbitrary group categories, such as race and gender (42). For instance, when someone says “Whites first” or “men first,” that person is rightly criticized. But “our nation first” is more widely acceptable, whether because national cooperation creates moral obligations, or because living under common laws, moving toward a common destiny, or sharing a common history establishes a morally justified preference for co-nationals. At the same time, unlike the family or the clan, a nation is an imagined and abstract community (43-44), and members are strangers to each other (45), in the same way that “citizens of the world” are strangers to each other. As a result, outgroup preference in the world-nation comparison may be more plausibly compared than the family-stranger comparison that some utilitarian thinkers examine.

Converging with recent evidence on the flexibility of local versus global moral judgments (46-48), our work offers global evidence of moral cosmopolitan judgment, as well as evidence of its dual factors, psychological sources, cognitive mechanisms, and associated interventions. In the context of the many intergroup moral problems the world faces - restrictions on refugees, climate change, pandemics, regulation of global trade, the public good of global scientific

collaboration - our research may offer some hope that the human capacity for reasoned judgment may be up to the challenge.

Methods

Scale. To develop and validate the new scale, we conducted six studies with 7388 total participants (SI1.1-1.4. Scale). In this iterative process, we designed and refined online surveys asking participants to indicate the degree to which they agree or disagree with each statement (from 1 *strongly disagree* to 6 *strongly agree*). Study 1.1-1.3 were exploratory tests of face validity and incorporated expert opinions. Study 1.4 and 1.5 were confirmatory tests on both American and Chinese samples, as well adding relevant psychological scales to test convergent and divergent validity, such as moral foundation scales and moral expansiveness scales. Study 1.6 was a full launch of our refined scale with international participants, which included behavioral measures of real-world charitable donations and free responses. These data appeared in our main analysis in the main text.

Predictive validity (SI2.1-2.5. Predictive analysis), in charitable donation, participants read:

We invite you to make a donation to a charitable cause. You can donate up to 10 of your Shells to one or both of the following organizations: Teach for All, an international NGO, or Teach for America [Teach First UK, and other country equivalent], a national NGO. The mission of both organizations is to improve education and expand opportunities for children in disadvantaged communities. Teach for All works worldwide, while Teach for America works in the US. Your donation decision will have a real-world effect, as we will donate as the way you suggested. Please indicate how much would you like to donate to: Teach for All (slider from 0 to 10) and Teach for American (slider from 0 to 10).

Their total must sum to 10. In the self-interest condition, participants read: “You could either keep the bonus to yourself, or donate up to 10 of the bonus to one or both of the following organizations.”

In the international trolley dilemma, participants read:

Suppose a runaway trolley is speeding down a track toward five middle-class citizens of our country. You are on a bridge over the tracks, and you can stop the trolley by putting something very heavy in front of it. As it happens, there is an immigrant next to you, and your only way to stop the trolley is to push him over the bridge and onto the track, killing him to save five. Is pushing him the right or wrong thing to do? 1 is *completely wrong* to 6 is *completely right*.

In the COVID petition (conducted in August, 2020), participants read:

In this section, we invite you to have your voice heard by the US Congress. On the next page, you will have a chance to influence a COVID-19 related petition we will send to Congressional Representatives on behalf of this study's participants. We will not mention your name, just the total number of participants and your collective choice. If you think that a petition should receive more support, please indicate by clicking on the retweet button. Retweet, broadly, means sharing the petition with other people. Just as in real life, not every retweet will get a response. You will choose one of the two petitions. We will send only one of these two petitions to members of Congress. The petition that receives most retweets will be sent out. So the more times you retweet, the more likely your desired petition, and the position you advocate, will be heard.

On the next page, they read:

People all over the world are interested in receiving vaccines against COVID-19. But people advocate for different strategies on how those vaccines should be made and distributed. Petition A supports Vaccinate [Mask] the World. It proposes that the US government and companies should make, sell, and donate COVID-19 vaccines [masks] to everyone in the world. Petition B supports Vaccinate [Mask] America. It proposes that the US government and companies should make, sell, and donate COVID-19 vaccines [masks] to our own citizens. On the next page, you will have a chance to retweet for the petition you prefer. They then look at two retweet buttons, under each petition. After clicking the retweet, a number will show up reminding participants how many times they have already clicked.

Regarding personal opinions on world policies, we prepared patriotism (e.g., I love my country), immigration (e.g., I would not like to have immigrants or foreign workers as my neighbors), loyal resistance (e.g., Sometimes a true patriot has to criticize our leaders), evaluations of international organizations (e.g., In general, international institutions do valuable work), and CO2 opinions (e.g., Our country should reduce its carbon footprint because climate change is causing some islands nations to disappear). (details in SI2.5. Policy attitudes)

Mechanism from linguistic data. For linguistic modeling (SI3.1-3.2. Natural sentences), with international samples, we google translated (API) all participants' responses into English and trained our computational linguistic models to classify moral cosmopolitans from moral patriots. Our model is based on Long-Short Term Memory Recurrent Neural Networks with self-attention 36,37) to predict donations from the text inputs. Just as humans do not process a full picture to understand what is going on but selectively attend to useful information, we trained our model to selectively detect influential linguistic cues by explicitly requiring the computation

of attention weights for each input word when making predictions. For example, our linguistic model predicts the following sentence, “both are important,” to be moral cosmopolitan, while also yielding the attention weights for each word (.42, .14, .23), which means that “both” receives the most weight (.42) and contributes significantly to this prediction.

For frequency counts (SI3.3. Single tokens), in American post-COVID samples, we gathered participants’ responses in terms of five single words. After making retweet decisions, participants read:

Please list briefly some thoughts you had while deciding and while retweeting. Single words are the best, one in each box. Please list at least one word, up to five.

We then classified these responses into Mask the World, Mask America, Vaccinate the World, and Vaccinate America groups. After basic data cleaning (e.g., remove stop words, lowercase, lemmatization), we counted top words in each category (word clouds in SI3.3).

Experiment. See pre-registration report for design and model details (SI4.1-4.2. Intervention). We present experimental stimuli here.

Cosmopolitan treatment condition:

(1) I think that anyone in the world has the right to education, it is something similar to food, so it must be donated in equal parts to both organizations. (2) I believe everyone should have the same opportunity to earn a quality education, both in the US and all over the world. I would like to see other countries flourish and improve. (3) There is an immediate need for education worldwide especially for women and children living in impoverished countries. (4) Worldwide many students especially GIRLS cannot get an education any other way... (5) Education is of paramount importance, both here and abroad. (6) I chose to donate half of the money to Teach for All because I think

international education needs more help. Though I am keenly aware of the educational issues in America, I think that the world problem is much larger and needs more help. (7)

I think providing education across world is important rather than just America. (8)

Because the first [organization] acts globally and that includes America.

Patriot treatment condition:

(1) I think we should help our own first then broaden the scope. (2) need to help our own country first. (3) We need to take care of our own people and stop help[ing] the world -

our country needs help for our children. (4) I believe we should help Americans first. (5)

While I would like to support children in other countries, domestic issues are a priority for me. (6) While education is needed throughout the world, we need to remember that

our USA children are in need also. (7) I think education opportunities should begin at

home in America. However, I do recognize the need for assistance in other countries as

well. (8) Education is good for everyone but more should be put at home.

/end

References

1. Ellemers, N. (2017). *Morality and the regulation of social behavior: Groups as moral anchors*. Routledge.
2. Greene, J. D. (2013). *Moral tribes: Emotion, reason, and the gap between us and them*. Penguin.
3. Janoff-Bulman, R., & Carnes, N. C. (2016). Morality: Hypo-egoic for us but not them. In K. Brown & M. Leary (Eds.), *Oxford handbook of hypo-egoic phenomena* (pp. 175-188). NY: Oxford University Press.
4. Choi, J. K., & Bowles, S. (2007). The coevolution of parochial altruism and war. *science*, 318(5850), 636-640.
5. Bruneau, E. G., Cikara, M., & Saxe, R. (2017). Parochial empathy predicts reduced altruism and the endorsement of passive harm. *Social psychological and personality science*, 8(8), 934-942.
6. De Dreu, C. K., Balliet, D., & Halevy, N. (2014). Parochial cooperation in humans: forms and functions of self-sacrifice in intergroup conflict. In *Advances in motivation science* (Vol. 1, pp. 1-47).
7. Hughes, B. L., Ambady, N., & Zaki, J. (2017). Trusting outgroup, but not ingroup members, requires control: neural and behavioral evidence. *Social cognitive and affective neuroscience*, 12(3), 372-381.
8. Molenberghs, P., Gapp, J., Wang, B., Louis, W. R., & Decety, J. (2014). Increased moral sensitivity for outgroup perpetrators harming ingroup members. *Cerebral cortex*, 26(1), 225-233.

9. Romano, A., Balliet, D., Yamagishi, T., & Liu, J. H. (2017). Parochial trust and cooperation across 17 societies. *Proceedings of the national academy of sciences*, 114(48), 12702-12707.
10. Davidai, S., & Ongis, M. (2019). The politics of zero-sum thinking: The relationship between political ideology and the belief that life is a zero-sum game. *Science advances*, 5(12), eaay3761.
11. Pogge, T. (2005). World poverty and human rights. *Ethics & international affairs*, 19(1), 1-7.
12. Singer, P. (1981). *The expanding circle*. Oxford: Clarendon Press.
13. Singer, P. (2016). *One world now: The ethics of globalization*. Yale University Press.
14. Cappelen, A., Gauri, V., & Tungooden, B. (2018). *Cooperation creates special moral obligations*. Cesifo Working Papers 7052.
15. Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in cognitive sciences*, 17(8), 413-425.
16. McFarland, S., Brown, D., & Webb, M. (2013). Identification with all humanity as a moral concept and psychological construct. *Current directions in psychological science*, 22(3), 194-198.
17. Crimston, D., Bain, P. G., Hornsey, M. J., & Bastian, B. (2016). Moral expansiveness: Examining variability in the extension of the moral world. *Journal of personality and social psychology*, 111(4), 636.
18. Graham, J., Waytz, A., Meindl, P., Iyer, R., & Young, L. (2017). Centripetal and centrifugal forces in the moral circle: Competing constraints on moral learning. *Cognition*, 167, 58-65.

19. Mummendey, A., & Otten, S. (1998). Positive–negative asymmetry in social discrimination. *European Review of Social Psychology*, 9(1), 107-143.
20. Allport, G. W. (1954). *The nature of prejudice*. Addison-Wesley.
21. Brewer, M. B. (1999). The psychology of prejudice: Ingroup love and outgroup hate? *Journal of social issues*, 55(3), 429-444.
22. De Dreu, C. K., Balliet, D., & Halevy, N. (2014). Parochial cooperation in humans: forms and functions of self-sacrifice in intergroup conflict. In *Advances in motivation science* (Vol. 1, pp. 1-47). Elsevier.
23. Halevy, N., Bornstein, G., & Sagiv, L. (2008). “In-group love” and “out-group hate” as motives for individual participation in intergroup conflict: A new game paradigm. *Psychological science*, 19(4), 405-411.
24. Cleveland, M., & Laroche, M. (2007). Acculturaton to the global consumer culture: Scale development and research paradigm. *Journal of business research*, 60(3), 249-259.
25. Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of personality and social psychology*, 101(2), 366.
26. Rosseel, Y. (2012). Lavaan: An R package for structural equation modeling and more. Version 0.5–12 (BETA). *Journal of statistical software*, 48(2), 1-36.
27. Higgins, E. T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. In *Advances in experimental social psychology* (Vol. 30, pp. 1-46). Academic Press.

28. Janoff-Bulman, R., Sheikh, S., & Baldacci, K. G. (2008). Mapping moral motives: Approach, avoidance, and political orientation. *Journal of experimental social psychology*, 44(4), 1091-1099.
29. Cikara, M., Farnsworth, R. A., Harris, L. T., & Fiske, S. T. (2010). On the wrong side of the trolley track: Neural correlates of relative social valuation. *Social cognitive and affective neuroscience*, 5(4), 404-413.
30. Foot, P. (1967). The problem of abortion and the doctrine of double effect.
31. Kahane, G., Everett, J. A., Earp, B. D., Caviola, L., Faber, N. S., Crockett, M. J., & Savulescu, J. (2017). Beyond sacrificial harm: A two-dimensional model of utilitarian psychology.
32. Hesse, M, Why might women be less eager to get the coronavirus vaccine? An investigation. *Washington Post*. https://www.washingtonpost.com/lifestyle/style/women-covid-vaccine-skeptical/2020/12/15/63551cac-3a61-11eb-9276-ae0ca72729be_story.html. Accessed 15 December 2020.
33. Ames, D. L., & Fiske, S. T. (2015). Perceived intent motivates people to magnify observed harms. *Proceedings of the National Academy of Sciences*, 112(12), 3599-3605.
34. Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of personality and social psychology*, 44(1), 113.
35. Ho, A. K., Sidanius, J., Kteily, N., Sheehy-Skeffington, J., Pratto, F., Henkel, K. E., ... & Stewart, A. L. (2015). The nature of social dominance orientation: Theorizing and

- measuring preferences for intergroup inequality using the new SDO₇ scale. *Journal of Personality and Social Psychology*, 109(6), 1003.
36. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9, 1735-1780.
37. Lin, Z., Feng, M., Santos, C. N. D., Yu, M., Xiang, B., Zhou, B., & Bengio, Y. (2017). A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*.
38. Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In *Communication and persuasion* (pp. 1-24). Springer, New York, NY.
39. Fisher, W. R. (1985). The narrative paradigm: An elaboration. *Communications monographs*, 52(4), 347-367.
40. Pinker, S. (2011). *The better angels of our nature. why violence has declined*. Viking.
41. Hoover, J., Johnson, K., Boghrati, R., Graham, J., & Dehghani, M. (2018). Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation. *Collabra: Psychology*, 4(1).
42. Rawls, J. (1993). The law of peoples. *Critical inquiry*, 20(1), 36-68.
43. Anderson, B. (2006). *Imagined communities: Reflections on the origin and spread of nationalism*. Verso Books.
44. James, P. (1992). Forms of abstract "community" from tribe and kingdom to nation and state. *Philosophy of the social sciences*, 22(3), 313-336.
45. Packard, V. O. (1972). *A nation of strangers*. McKay.

46. Bian, L., Sloane, S., & Baillargeon, R. (2018). Infants expect ingroup support to override fairness when resources are limited. *Proceedings of the national academy of sciences*, *115*(11), 2705-2710.
47. Gaesser, B., Shimura, Y., & Cikara, M. (2019). Episodic simulation reduces intergroup bias in prosocial intentions and behavior. *Journal of personality and social psychology*.
48. Kleiman-Weiner, M., Saxe, R., & Tenenbaum, J. B. (2017). Learning a commonsense moral theory. *Cognition*, *167*, 107-123.

Tables.

Table 1. Final scale items for moral cosmopolitanism

Subscale	No.	Item
Equitable Benefit (EB)	1	Citizens in our country should accept a reduction in their standard of living if doing so increases global equality.
	2	Citizens in our country are obliged to lower their standard of living if doing so substantially contributes to global justice.
	3	Citizens in our country must lower their standard of living if that is necessary to achieve a higher standard of living for every person in the world.
	4	It is a problem if most people in our country generally agree on the principle of putting our own citizens first.
	5	It is wrong to always prioritize citizens over people in other countries.
Equitable Security (ES)	6	It is as bad when terrorist attacks happen elsewhere as when they happen in our nation.
	7	When a hostage is being held by terrorists, it is equally bad if the hostage is from other nations and from our nation.
	8	It is as bad when our military accidentally kills foreign civilians as when a foreign military accidentally kills our civilians.
	9	When our friend, who is a fellow-citizen, makes a distasteful joke about others, it is equally inappropriate when the joke concerns foreigners as when it concerns fellow citizens.

Figures.

Fig 1. Self-reported moral cosmopolitanism, Equitable Benefit (green) and Equitable Security (red), by country means and 95% confidence interval, raw scores scattered in the background. Country names on the y-axis, with geographical locations in the world map: United States (US), Guatemala (GT), Colombia (CO), Spain (ES), United Kingdom (UK), Germany (DE), China (CN), Japan (JP). On the x-axis, scores approaching 6 indicate relatively moral-cosmopolitan judgment, while scores approaching 1 indicate relatively moral-patriotic judgments.

Fig 2a.

Fig 2b.

Fig 2. Relations of cosmopolitan moral judgments with (a) comparative charitable donations and (b) international trolley dilemma, displaying country central tendency and 95% confidence intervals. Full model estimates individual-level linear effects while controlling for within-country dependencies with random slopes and random intercepts estimation. X-axis denotes self-reported cosmopolitan moral judgments. Y-axis denotes donation proportion (2a), a higher value indicates more to Teach for All; trolley dilemma decisions (2b), a lower value indicates reluctance to push an immigrant to save five citizens.

\$ to Teach for All	Responses to Why Donated
\$0	i believe in taking care of our own citizen before the rest of the world
\$0.1	charity begin at hom
\$0.2	charity start at home so first help those around you then yourself then everyone you can
\$0.3	should put our own house in order first
\$0.4	because it is important that the educational standard in uk are maintained and education should be a priority to improve the educational standard of the world
\$0.5	because both have the same miss
\$0.6	the world project reach more peopl
\$0.7	i think providing education across world is important rather than just america
\$0.8	i see humanity a a whole and not divided into country so the bulk of the donation should go to a global organization nevertheless i would also donate a small amount locally so that the level of education in germany increase and it ha a direct influence on u
\$0.9	international research is more import
\$1	all place in the world can benefit

Fig 3. Attention heatmap learned from linguistic model using google-translated responses from eight countries. See SI 3.2 for model architecture details. Each line is an example response from patriotic (\$0 to Teach for All) to cosmopolitan (\$1) donations. Red denotes model estimated word weights, the darker the more important a word is.

Fig 4. Average treatment effects on international donation by American online participants. Null hypothesis of no difference in effect was rejected, by an 11.8% increase. Plot features central tendency and kernel density estimations for experimental conditions.