

Trade-offs in Economic Preferences: Both Democrats and Republicans Prioritize Poverty Alleviation over Equality and Efficiency

Dylan Wiwad¹, Azim Shariff², & Nour S. Kteily¹

¹Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL, 60208, USA.

²Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, BC, V6T 1Z4

Corresponding Authors:

Dylan Wiwad

Management and Organizations

Kellogg School of Management, Northwestern University

2211 Campus Drive

Evanston, IL, 60208, USA

Email: dylan.wiwad@kellogg.northwestern.edu

Nour S. Kteily

Management and Organizations

Kellogg School of Management, Northwestern University

2211 Campus Drive

Evanston, IL, 60208, USA

Email: n-kteily@kellogg.northwestern.edu

Author Contributions: D.W, A.F, and N.S.K designed research, D.W and N.S.K performed research, D.W and N.S.K analyzed data, and D.W, A.F, and N.S.K wrote the paper.

Competing Interest Statement: Disclose any competing interests here.

Classification: Social Sciences (Major) and Psychological and Cognitive Sciences (Minor)

Keywords: Paste the keywords here. There should be at least three and no more than five.

This PDF file includes:

Main Text

Figs. 1 to 3

Table 1

Abstract

Tackling socioeconomic woes can be a difficult task in an increasingly polarized political climate, especially when desired outcomes conflict (e.g., alleviating poverty but increasing inequality). Across two main (N = 1,931) studies and one supplemental (N = 956) study we develop a typology that orthogonalizes social policy along two dimensions: (a) help versus harm and (b) the rich versus the poor. In doing so, we explore partisan overlap in the prioritization of poverty, inequality, economic efficiency, and wealth in determining political policy support. In Study 1 we find a large degree of partisan overlap; Republicans and Democrats alike strongly prioritize helping the poor and share a general dislike for helping the rich. Once the bipartisan goal of helping the poor is met, Democrats more strongly prioritize inequality reduction while Republicans more strongly prioritize avoiding harm to the rich. Additionally, Democrats are more willing than Republicans to forego help to the poor if it comes along with rising inequality. In Study 2 we replicate these findings and explore the emotional underpinnings of socioeconomic preferences. While only hostile emotions towards the rich predict Republicans' willingness to harm the rich and reduce inequality, both compassion towards the poor and hostile emotions towards the rich predict this preference among Democrats. By leveraging our comprehensive typology framework, we can generate valuable insights for designing social policy that falls at the intersection of bipartisan economic priorities.

Significance Statement

When designing social policy, scientists and policymakers alike debate whether poverty or inequality is the more pressing concern. We develop a typology to test how concerns about inequality, economic efficiency, poverty, and wealth independently, and jointly, influence policy support. Our results suggest that, contrary to previous research identifying singular motives, socioeconomic preferences reflect a complex interaction of competing concerns. While we find a large degree of partisan overlap in policy support, particularly on the priority of poverty reduction, we find some key differences in what Republican and Democrats are willing to tolerate to reduce inequality. These findings shed light on the psychology of socioeconomic decision making and provide insights for shaping political policies aimed at garnering broad bipartisan agreement.

Introduction

Economic inequality—the gap between the rich and the poor—is gaining societal and scientific attention (e.g., 1, 2). Some argue this attention is overdue, citing manifestations of economic inequality such as the \$2.6 trillion drop in the United States' GDP (3, 4) and the wealth transfer of nearly \$50 trillion since World War II from working class Americans to the top 1% (5). Further, over the last two decades, growing inequality within countries has been linked to a plethora of social ills, including worse public health, more violent crime, and increased psychological distress (1, 2, 6–9). Others, however, argue that the focus on inequality is misguided, suggesting that inequality per se is not the problem. Rather, we should be more specifically concerned about poverty—lacking the financial resources to meet one's basic needs—which causes hunger, sickness, and impaired decision-making (10–13). An excessive focus on inequality reduction, some argue, causes us to eschew policies that would raise the standard of living for *everyone*—including the poor—simply because they sometimes help the rich even more than the poor.

Here, we explore what drives individuals' social policy preferences generally, and in particular as they manage trade-offs between inequality and poverty alleviation. We decompose social policy into its constituent parts—constructing a 2x2 typology of policies that, on one dimension, help or harm, and on another dimension, affect the rich or poor—to build a more nuanced understanding of the motivations underlying individuals' policy support. Are individuals willing to forgo help to the poor in order to constrain rising inequality, or does helping the poor take precedence? What emotions underlie these preferences, and how do individuals' ideological leanings shape their perspectives?

The answers to these questions have important implications for how policy makers design solutions to contemporary socioeconomic problems. Consider debates about Coronavirus vaccine distribution. On the one hand, some individuals have advocated distributing vaccine doses whenever and wherever they are available, to whomever is ready and willing to be vaccinated (e.g., 14). From this perspective, the singular priority ought to be getting as many total people vaccinated as quickly as possible. Such a policy would, proponents argue, maximize the welfare of everyone, including the poor—even if it risks exacerbating existing disparities, due to accessibility advantages in high-income and/or majority-white areas. On the other hand, for those who prioritize equality, it might be preferable to slow distribution until guardrails are in place to assure equality in distribution, even if that might mean that everyone—including the poor—gets vaccinated more slowly than they might otherwise have been (15).

Individuals might be drawn to one social policy or another not only as a function of its impact on the poor, its impact on the rich, and its impact on levels of inequality, but also as a function of concerns about economic efficiency (which we define here as the sum total outcome of a decision, distribution, or social policy on the system as a whole (16)). That is, some individuals might favor a policy that embraces first-come-first-served vaccination simply because these individuals prioritize maximizing the overall number of vaccines administered—irrespective of the policy's effects on the rich, the poor, or inequality.

Teasing apart the true reasons underlying individuals' policy preferences is challenging, because the various relevant dimensions are frequently confounded. For example, individuals who support a wealth tax could be specifically driven by a desire to harm the rich, a desire to help the poor, a desire to reduce inequality, or some combination thereof. Moreover, there is also the possibility of complex interactions between the factors driving policy support. For example, individuals might be willing to forgo some amount of potential gain to the poor if doing so stifles rising inequality. Additionally, there could be differences between gains and losses: for example, individuals might be willing to forgo gains to the poor in order to reduce inequality by a given amount but unwilling to impose an equivalent level of *costs* on the poor to reduce inequality by the same amount.

Despite the difficulty, teasing apart underlying socioeconomic preferences is critical for understanding and addressing partisan conflict in the United States. Many researchers have documented not only widening political polarization (17) but also rising interpersonal animosity towards the political outgroup (18). To what degree does this strife reflect fundamental disagreements on the most desired socioeconomic outcomes? For example, Democrats are sometimes seen as motivated by harming the wealthy as an end in and of itself. Conversely, Republicans are sometimes seen as motivated uniquely by boosting the wealth of the rich. When we strip away the ideologically loaded political context (e.g., tax breaks or expanding welfare) and instead focus on the root economic principles (i.e., the wealth of the rich and the poor, the level of inequality, economic efficiency) do partisans display more or less attitudinal overlap? In the present work we leverage our typology to more deeply explore this question.

While previous work has been sensitive to exploring how people prioritize poverty, inequality, and efficiency, this work has been hampered by important limitations. In a representative experimental design (e.g., 17–20), participants evaluate a set of three cash bonus payment distributions in lab-based choice tasks: one that is efficient, one that minimizes inequality, and one that maximizes the payout of the “poorest” player. Work investigating which of these factors people prioritize has resulted in conflicting findings, with some work supporting each of the three possibilities—concern for increasing efficiency (16, 21), minimizing inequality, (19–21), and helping the poor (22)—as the highest priority. Critically, the small number of distributions used in previous paradigms do not allow researchers to explore, across a large number of potential distributions, how one priority (e.g., inequality aversion) may become more or less important at differing levels of another (e.g., amount of help to the poor). Additionally, when individuals are picking between options in a given distribution set, choosing the inequality-minimizing option implies that the decider cares about reducing inequality and *not* efficiency or maximizing the welfare of the poor. Finally, previous work does not disentangle what is happening to the richest player from what is happening to the poorest player; it is impossible to know whether those who select the inequality-minimizing outcome are doing so out of a desire to restrict the gains of the rich or lessen the relative disadvantage to the poor (or both).

In principle, and as noted above, individuals’ preferences and priorities might be driven by complex considerations. That is, individuals’ preferences might factor in *all* of efficiency, inequality, and the welfare of the poor, with different dimensions emerging as more or less important at various levels of the other factors. Without considering a very large set of potential distributions—including cases where effects on the rich and poor, as well as inequality and efficiency, are orthogonal to one another—the types of conclusions that can be drawn about how individuals trade off these factors are necessarily limited.

In order to address these questions and better elucidate the various factors that interactively shape individuals’ social policy support, we created a typology that orthogonalizes the impact that policies can have along two dimensions: (a) help versus harm and (b) the rich versus the poor (Fig. 1a). In doing so, we can observe not only how changes to the outcome of the rich and the poor influence policy support, but also how changes to inequality and efficiency influence behavior. Particularly, we can empirically distinguish how strongly a person is independently motivated by each of these four factors.

Study Overview

In order to tease apart the effects of these complex and interconnected outcomes (and how they may differ along ideological lines), in Study 1, we introduce a typology that contains 440 policies spread across four quadrants (or, zooming in further, eight octants). The four quadrants are generated from two orthogonal axes that respectively represent policies encompassing a large degree of positive and negative impact to the wealth of (a) the rich and (b) the poor (Fig. 1a). Here, the policies ranged from a decrease of 50% to an increase of 50% (in 5% increments) in wealth for each group. We devised this structure in order to compare support for policies as a

function—and at the intersections—of four broad dimensions: (1) the degree to which these policies “help” (i.e., increase the wealth; each of the poor and rich), (2) the degree to which they “harm” (i.e., decrease the wealth; each of the poor and rich), (3) their impact on inequality (i.e., the wealth gap between the poor and rich), and (4) their impact on economic efficiency (i.e., the size of the overall economic pie, independent of how it is allocated). Note that all policies to the left of the y-axis decrease the wealth of (i.e., “harm”) the poor, whereas all policies to the right of the y-axis increase the wealth of (i.e., “help”) the poor.

All policies above the x-axis help the rich whereas all policies below the x-axis harm the rich. All policies above the positive sloping diagonal increase inequality (i.e., increasing the gap in wealth between the rich and poor) and all policies below the positive sloping diagonal decrease inequality (i.e., decreasing the gap in wealth between the rich and the poor). And all policies above the negative sloping diagonal result in an overall increase in wealth (i.e., result in a “net gain,” in percentage terms) whereas all policies below the negative sloping diagonal result in an overall decrease in wealth (i.e., result in a “net loss,” in percentage terms). This typology allows us to ask questions like “How much growth in inequality are people willing to tolerate in order to ensure gains to the poor?” and “How much harm to the rich will people tolerate in order to reduce inequality?”

We chose to use percentages in the typology because, while imperfect, they are easy to understand. In this paradigm, we treat a policy that helps the poor by 10% and helps the rich by 5% as reducing inequality (i.e., the gap in wealth is closing). We recognize that it is possible that a smaller percentage change of a larger overall sum (i.e., a 5% increase to the wealth of the rich versus a 10% increase to the wealth of the poor) might still increase inequality due to the vast difference in absolute resources. To address this disadvantage, we replicated Study 1 using raw dollar amounts—with policies ranging from a decrease of \$500 million to an increase of \$500 million in the wealth of the rich and the poor—and find a consistent overall pattern of results (See SI Study 1 Replication for full analyses).

Study 1 considers the role of political ideology in this process, an overarching factor that may be associated with whether individuals respectively prioritize factors such as efficiency, inequality, and poverty alleviation. In particular, individuals on the political left tend to rely more strongly on considerations of harm reduction and fairness in their moral reasoning (23) while those on the political right tend to be more tolerant of social and economic hierarchy (24). The political left’s prioritization of harm reduction and fairness may manifest primarily as support for those at the bottom of the social and economic hierarchy; relative to those on the right, individuals on the political left tend to show more empathy for (25), attention to (26), and efforts to overturn hardships (27) experienced by socially disadvantaged (vs. advantaged) groups, effects thought to be driven by a desire to attenuate inequality. Thus, we might expect that relative to those on the political right, individuals on the political left will (a) more strongly desire reducing harm (and providing help) to the poor; (b) be more averse to helping (and more tolerant of harming) the rich; and (c) be more opposed to increasing (and committed to decreasing) inequality. Additionally, based on prior research, we might expect those on the political right (vs. left) will more strongly desire efficiency (and more strongly oppose inefficiency; 26).

In Study 2, we replicate the results of Study 1 in a new sample, and further consider the role of emotional underpinnings in driving support for policies occupying various positions in our typology. For example, we ask why certain individuals prioritize policies that, while providing equivalent gains to the poor, either help or hurt the rich. Note that when holding effects on the poor constant, choosing to hurt versus help the rich has the corollary effect of reducing inequality. On the one hand, choosing to hurt the rich in this scenario might be driven by moral outrage at inequality (29, 30) and/or disdain towards the rich (31). Alternatively, this same preference could be driven by compassion or empathy for the *poor*—a tacit belief that inequality is psychologically bad for those in poverty and thus alleviating inequality will help them. Finally, it could be the case

that both hostile emotions towards the rich *and* compassionate emotions towards the poor drive such preferences. We consider these and related questions in our second sample.

Results

Study 1

We recruited a sample of 1,457 American adults on Amazon's Mechanical Turk to vote ("Would you vote for this policy?" Yes = 1, No = 0) on 30 (out of 440) randomly drawn hypothetical economic policies.¹ Our dataset is therefore Massively Missing Completely at Random (MMCAR, see Methods; (32)). The resulting dataset contained 43,710 individual votes spread approximately evenly across all policies (*mean votes per policy* = 102.74, *range* = 95-109). In order to visualize these votes within the typology, we created heatmaps of policy support from the raw data, broken down for the sample as a whole as well as for Democrats and Republicans separately (Figs. 1b-e).

New Typology Uncovers Trade-offs

We find that only one axis cleanly divides overall policy support: helping versus harming the poor. In particular, we see that policy support is extremely low for policies that harm the poor (Fig. 1b; $M_{support} = 13.33\%$), irrespective of their effects on the rich, inequality, or efficiency (Range in average support across octants that involve harming the poor = 12.29% - 14.39%). But, when examining policies that help the poor, support becomes more contingent on other factors. People broadly support policies that help the poor more than the rich—thereby reducing inequality ($M_{support} = 87.24\%$; Octant 2). However, support for policies that help the poor diminishes when the rich are helped more than the poor—thereby growing inequality ($M_{support} = 66.85\%$; Octant 1; $\chi^2(1) = 535.19, p < .001, 95\% \text{ CI } [.668, .874]$)—or the rich are harmed instead of helped (despite reductions in inequality; $M_{support} = 78.05\%$; Octant 3; $\chi^2(1) = 134.74, p < .001, 95\% \text{ CI } [.874, .781]$).

Although informative, descriptive explorations of heatmaps provide only a coarse window into the reasons why policy support changes as we cross axes of the typology. We cannot tell, for example, if the reduced support for policies that harm instead of help the rich when the poor are being helped is due to an aversion to harming the rich, a relative insensitivity to inequality, or some combination thereof. In order to isolate the effect of each dimension of a policy while holding the other dimensions constant we need to model how support for policies varies across the typology as a whole.

Helping vs. hurting the poor is the most important, but not only, predictor of policy support

In order to isolate the effect of each policy dimension, we conducted a set of standardized binomial crossed-factors multilevel regressions predicting policy support (1 = "Yes" vote, 0 = "No" vote) across our entire typology. We coded each policy within the typology according to whether it decreased (-1), did not change (0), or increased (+1) each of the wealth of the poor, the wealth of the rich, inequality, and economic efficiency (see Methods for full model descriptions and SI for model choice justification). In these models we nested the individual policy

¹ Additionally, participants reported their support for each policy on a 1 (Would Strongly Oppose) to 7 (Would Strongly Support) scale. While we do not report results using the Likert scale in this manuscript, it is worth noting that when conducting regression analyses with these items, all coefficients are significant and in the same direction. Data and analyses can be found on the OSF at https://osf.io/rt7uq/?view_only=8d2ae3bc4cc340a9bd525113657a59d4.

votes (n = 43,710; Level 1) simultaneously under both participants (n = 1,457; Level 2) as well as policies (n = 440; Level 2).

We found that, controlling for all other outcomes, the primary predictor of policy support was helping the poor (Table 1; Model 1a; all noted effects here are statistically significant). Helping (versus harming) the poor increased model-predicted policy support from 10% to 76% (Fig. S1a, green line). Reducing inequality (versus increasing inequality) was the second strongest predictor of policy support, increasing predicted policy support from 27% to 50% holding other factors constant (Fig. S1b, green line). Finally, participants were, overall, more likely to favor policies that were economically efficient and that helped the rich, but these effects were relatively weak. Holding all else constant, increasing (versus decreasing) economic efficiency bolstered predicted policy support from 30% to 45% (Fig. S1d, green line), while helping (versus harming) the rich bolstered predicted policy support from 33% to 43% (Fig. S1c, green line). See Fig. 2 and Table S1 for statistical comparisons between these four main effects and Table S2 for analyses controlling for age, income, and dummy coded gender and race.

Partisan Differences. We next considered this same analysis breaking down our data based on participants' political ideology. The partisan heatmaps (Figs. 1c-d) demonstrate clear patterns of convergence and divergence among Democrats and Republicans (See Study 1 Supplemental Results for a comparison of structural similarity across partisan heatmaps). Differences were most pronounced in the degree to which Democrats and Republicans (a) wish to help versus harm the rich and (b) prioritize inequality reduction versus expansion (Fig. 2; Table 1). Using the same crossed-factors multilevel regressions, we find that helping (versus harming) the poor is the primary concern across partisan lines (Democrats: +72 percentage points; Republicans: +59 percentage points), albeit slightly stronger amongst Democrats (Table 1 two-way interactions). Democrats (+27 percentage points) also place greater importance than Republicans (+18 percentage points) on inequality reduction (vs. expansion). Additionally, Republicans (+24 percentage points) place greater importance than Democrats (-4 percentage points) on helping (versus harming) the rich. Lastly, partisans do not differ in how strongly they weight efficiency (versus inefficiency; Democrats: +14 percentage points; Republicans: +16 percentage points).

Differentiating Help from Harm

Despite these insights, the previous analyses are limited by the fact that the coding scheme allows each hypothetical policy to receive only a categorical score of +1, -1, or 0 representing whether it corresponds to an increase, decrease, or no change on our dimensions of interest. Despite having several advantages (see SI for Model Choice Justification), this coding scheme overlooks meaningful variation in the *strength* of endorsement within those categories; for example, policies coded simply as “1 – helps the poor” vary substantially in how *much* of an increase to the poor they offer (i.e., from +5% to +50%). Additionally, whereas the previous model treats helping and harming as opposite ends of a single continuum (i.e., -1 to 0 to +1), people may react differently to help versus equivalent harm (e.g., individuals' aversion to harming the rich might be stronger than their corresponding preference for helping them).

Thus, in order to differentiate help from harm, we conducted an additional analysis that takes *degree* of change into account and treats increases and decreases as separate dimensions. While we broadly replicate the findings of the earlier analyses (See SI Study 1 Results and Table S3 for full analyses of this model; see Table S4 demonstrating that all analyses comparing Republicans and Democrats hold when controlling for age, income, and gender), this model provides additional nuance to our understanding of how concerns about the rich motivate policy support across ideological lines. Broadly, whereas we previously observed no unique effect for Democrats of harming versus helping the rich on policy support (suggesting an indifference to the wealth of the rich; Fig. 3), we find in this model that Democrats actively oppose *both* large degrees of help and large degrees of harm to the rich (Table S3). That is, while

Democrats prefer to not give additional help to the rich, they also prefer not to levy a great deal of harm on to the rich. Similarly, we find that Republicans are strongly averse to large degrees of harm to the rich but are actually also weakly averse to large degrees of help to the rich (Table S3). Thus, Republicans are more concerned about avoiding harm than providing gains to the rich, and in fact share (albeit more mildly) Democrats' distaste for large gains to the rich. Moreover, while Democrats oppose granting the rich more gains, they nevertheless also bid them no harm.

In Making Trade-offs to Help the Poor, Republicans Tolerate Greater Increases to Inequality than Democrats

Thus far we were able to explore how various factors (e.g., helping versus harming the poor and rich, increasing versus decreasing inequality and efficiency) independently drive policy support across a complex array of interdependent motives, all else held equal. However, in both the typology and practice, many of the situations we face (and economic policies we must choose between) require making difficult trade-offs across various dimensions we might care about. Indeed, it is these very conflicts that lie at the heart of differences in opinion and interpretation about social trends (e.g., increasing globalization) over the past century, in which standards of living have risen for all but much more so for the rich than the poor (10, 11). How do individuals manage such trade-offs? Here, we focus on one trade-off in particular: What is the maximum amount of increased inequality people will countenance in order to help the poor?

To test this question, we regressed policy support on to the degree to which a policy (a) helps the poor, (b) increases inequality, as well as political party identification (see Table S5). That is, we restricted the data to only the cases where the poor are helped and inequality increases, “zooming in” on policies where these outcomes are in conflict (Fig. 1a, Octant 1). Consistent with our prior analyses, we see that the more the poor are being helped ($\beta = 1.077$, $p < .001$, 95% CI [.815, 1.339]) and the less inequality is increasing ($\beta = -1.793$, $p < .001$, 95% CI [-2.144, -1.431]) the higher the policy support. Critically, we can compute the interaction between helping the poor and increasing inequality, finding that these factors do indeed depend on one another ($\beta = 1.414$, $p < .001$, 95% CI [.835, 1.989]; Table S5). This interaction suggests that how strongly people prioritize preventing growth in inequality depends on how much the poor are being helped.

To examine how individuals manage trade-offs between increasing inequality and helping the poor, we identified the point of increased inequality at which support for a policy would drop below 50% at varying levels of help to the poor. Analyses revealed that both Democrats and Republicans tolerate a relatively large degree of increase to inequality as long as the poor are being helped even a small amount (Fig. 3). For example, our model predicts a 73% chance that a Democrat and a 92% chance that a Republican supports a policy that increases inequality by 15% so long as it helps the poor by just 5% (Fig. 3).² However, there are limits to individuals' tolerance of increasing inequality, even when this means forgoing gains to the poor. This model predicts that, for example, Democrats would fall below likely support (50%) for policies that help the poor by 5% if it comes along with a 30% increase in inequality (Fig. 3). Conversely, at a 5% gain to the poor Republicans would only fall below majority policy support if inequality increases by 45% or more (Fig. 3). Thus, Republicans appear to tolerate higher increases to inequality than

² Note that a limitation of this analysis is the fact that the ceiling on the change to inequality in our analyses (+50%) becomes smaller as a policy increases in the extent to which it helps the poor. For example, a policy that helps the poor by 40% can, in our paradigm, only increase inequality by a maximum of 10% (because the rich can, at most, be helped 50%, i.e., 10% more than the poor). Because of this, the model must increasingly extrapolate beyond the data as help to the poor increases. Crucially, this is true across ideological lines, and therefore any cross-partisan comparisons are unaffected. Regardless, this analysis is *most* informative when you consider support for policies that help the poor more modestly.

Democrats provided the poor are helped to some degree. Put differently, Democrats are faster than Republicans to forgo (some) gains to the poor if doing so stifles rising inequality.

Summary

Across a series of heatmaps and multilevel models we found consistent evidence of a broad convergence between the priorities of Democrats and Republicans. Indeed, across partisan lines, helping (versus harming) the poor was the most important outcome in determining policy support, and both Democrats and Republicans generally favor reducing (versus increasing) inequality and increasing (versus reducing) economic efficiency. Indeed, Democrats and Republicans appear to agree on more than they disagree with respect to the drivers of their economic policy support. Nevertheless, Democrats and Republicans do diverge meaningfully in the extent to which they prioritize reducing inequality and avoiding harm to the rich, as well as how they manage the trade-off between helping the poor and increasing inequality. Republicans are much more averse to harming (and supportive of helping) the rich than Democrats are. Additionally, Democrats put a higher priority on reducing (versus increasing) inequality than Republicans do. When desired outcomes come into conflict (e.g., helping the poor but increasing inequality) we found that Democrats are faster than Republicans to forgo gains to the poor if doing so prevents rising inequality. Conversely, in a supplemental analysis (See SI Supplemental results) we found that Democrats will tolerate policies that levy up to a 50% wealth reduction on the rich so long as the policy helps the poor and reduces inequality. Republicans, however, were unwilling to tolerate policies that impose large degrees of harm to the rich (i.e., greater than 35%) irrespective of how much these policies reduce inequality and help the poor.

Study 2

We had three aims in Study 2. First, we sought to replicate the results from Study 1 in a new sample. Second, we briefly explored the external validity of the typology by considering how policy support is related to real-world voting for political candidates that shared relevant features (i.e., Bernie Sanders versus Joe Biden). Finally, we sought to examine the emotional underpinnings of policy support within the typology.

We recruited a sample of 474 American adults on Amazon's Mechanical Turk to vote ("Would you vote for this policy?" Yes = 1, No = 0) on a full set of 119 hypothetical political policies. The procedure was nearly identical to Study 1, except that we utilized a complete observation design rather than a MMCAR design (see Methods). In order to facilitate this design and avoid participant fatigue, we made two changes to the scope of the typology. First, we restricted the policies to range from plus or minus 35% (as opposed to 50%). Second, we removed the policies that hurt the poor (given the extremely low bipartisan support for such policies). Thus, the resulting dataset contained 56,402 individual votes spread evenly across all policies (474 votes per policy). Additionally, participants reported compassion as well as negative (i.e., anger and disgust) emotions towards both the rich and poor.

Conceptual Replication of Study 1

We first replicated the key analyses from Study 1 (Model 1). Consistent with Study 1, all predictors of policy support were significant and in the same direction (Table S6). That is, we again found that controlling for political ideology, helping (vs. not affecting) the poor ($\beta = 0.484$, $p < .001$, 95% CI [.425, .542]), helping (vs. harming) the rich ($\beta = 0.290$, $p < .001$, 95% CI [.202, .377]), reducing (vs. growing) inequality ($\beta = -0.798$, $p < .001$, 95% CI [-.872, -.723]), and increasing (vs. decreasing) efficiency ($\beta = 0.436$, $p < .001$, 95% CI [.362, .510]) all predicted greater policy support. We also replicated the analyses taking degree of policy support into account (Fig. S2). Although we broadly replicated our main conclusions, there were some subtle shifts regarding the rank ordering and partisan differences (See SI Supplemental Results for full exploration of these differences).

Typology Voting Patterns Among Democrats Differentiate Biden Supporters from Sanders Supporters

In order to test whether and how preferences expressed in the paradigm relate to real-world political preferences, we examined how well voting within the typology mapped on to reported voting for Joe Biden or Bernie Sanders in the 2020 Democratic Primaries among our sample's 182 self-reported registered Democrats who voted for either candidate ($n_{votes} = 21,657$). At the time, Biden and Sanders represented competing candidates and different ideological subsets within the Democratic party (with Donald Trump as their incumbent candidate, no comparable primary competition occurred that year among the Republican party). For each Democrat-identifying participant we computed weights reflecting the proportion of policies they voted yes on within a given octant. For example, if a participant voted "yes" on 16 out of 21 policies in Octant 2, their weight for that octant would be .762.

We found meaningful differentiation between Biden and Sanders voters. In particular, we found that controlling for support for policies in each of the other octants, only stronger support for policies that harm the rich more than they help the poor (thereby reducing inequality; Octant 4 in Fig. 1a) predicted a greater likelihood of voting for Bernie Sanders over Joe Biden ($\beta = 0.621$, $p = .029$, 95% CI [.063, 1.178]; See Table S7 for full model). Beyond this, we predicted typology policy support among Biden and Sanders voters as a function of change to the poor, rich, inequality, and efficiency (-1 = decrease, 0 = no change, 1 = increase; Table S8). Consistent with themes from Sanders' and Biden's political platform—and in particular Sanders' emphasis on the objectionability of rising wealth at the top versus the bottom—we found that harming (versus helping) the rich predicts higher policy support among Sanders supporters (+16 percentage points) than Biden supporters (for whom, similar to Republicans, harming versus helping the rich actually *reduced* policy support; -5 percentage points). Additionally, whereas helping (versus harming) the poor was a driver of policy support for both Sanders and Biden supporters, it was a stronger predictor among Sanders supporters (Sanders: +53 percentage points; Biden: +39 percentage points). Interestingly, Biden (+42 percentage points) and Sanders (+42 percentage points) supporters placed a similar emphasis on inequality reduction (vs. expansion). Lastly, Biden supporters (+27 percentage points), relative to Sanders supporters (+19 percentage points), weight efficiency (versus inefficiency) more heavily when voting within the typology.

These findings suggest that real-world support for Democratic candidate Bernie Sanders over Joe Biden may be less about differential concern for inequality per se than by relatively greater willingness to harm (vs. help) the rich, greater concern for the poor, and lower concern for economic inefficiency. When we further compare across ideological lines, we find that Biden supporters fall decidedly in between Republican and Sanders voters (See Figs. S3a-c). These analyses demonstrate that this typology can find meaningful differentiation and predict important outcomes between ideologically similar people (e.g., within Democrats), even in contexts as important and impactful as real-world voting behavior. In our final analyses, we turn to testing the broader emotional underpinnings of policy support to get a deeper understanding of what is driving political voting behavior across partisan lines.

Positive Emotions Towards the Poor are Just as Important as Negative Emotions Towards the Rich in Predicting Willingness to Harm the Rich to Reduce Inequality

We have seen thus far that people, especially Democrats, support policies that harm the rich (Fig. 1e; Octants 3 and 4)—why? One proposition is that individuals are motivated to support policies that harm the rich (even when the poor are unaffected) simply due to anti-rich sentiment (i.e., anger and disgust towards the rich). Another is that support for such policies is in fact motivated, at least to some extent, by compassion towards the *poor*. In particular, individuals might hold that *inequality* per se imposes costs on the poor and thus that harming the rich will—by virtue of reducing inequality—increase the well-being of the poor.

In order to test this possibility, we explored the relative preference for policies that harm versus help the rich while holding help to the poor constant. On average, Quadrant 1 (comprising Octants 1 and 2) and Quadrant 2 (comprising Octants 3 and 4; Fig. 1a) offer the same degree of help to the poor. They differ, however, in whether they help the rich and have a net neutral effect on inequality (Quadrant 1) or harm the rich and reduce inequality (Quadrant 2). Because help to the poor is constant, one might expect that compassion towards the poor would be irrelevant in predicting support for one quadrant over the other, and that relative support for Quadrant 2 over Quadrant 1 (i.e., harming versus helping the rich and thereby reducing inequality) would be driven solely by hostility towards the rich. On the other hand, individuals might believe that inequality itself imposes psychological costs on the poor, and therefore that reducing inequality would be beneficial to the poor even if it solely involved taking things away from the rich. From that perspective, the relative preference for harming the rich might be predicted (at least in part) by compassion towards the *poor* even when these policies involve the same degree of benefit to the poor. These two possibilities paint substantially different pictures regarding the underlying motivations of those who advocate for inequality reduction.

Among Democrats, the relative preference for policies in Quadrant 2 versus Quadrant 1 (i.e., policies that hurt versus harm the rich while holding gains to the poor constant) is uniquely associated both with compassion towards the poor ($\beta = 0.350, p < .001, 95\% \text{ CI } [.191, .509]$) and hostility towards the rich ($\beta = -0.375, p < .001, 95\% \text{ CI } [-.516, -.234]$). Of note, these two emotional motivations among Democrats were equal in strength (Estimate of magnitude of difference = 0.025, $t(247) = 0.206, p = .837$). Conversely, among Republicans this relative preference was only associated with negative emotions towards the rich ($\beta = -0.426, p < .001, 95\% \text{ CI } [-.644, -.207]$), and not positive emotions towards the poor ($\beta = 0.110, p = .251, 95\% \text{ CI } [-.077, .300]$; Estimate of the magnitude of difference = 0.316, $t(220) = 2.355, p = .020$). These findings suggest that Democrats believe that inequality itself is detrimental to the poor and that reducing inequality will help them.

Discussion

The political climate in the United States continues to become increasingly polarized (18). This environment can make it challenging to introduce legislation tackling important socioeconomic woes such as economic inequality and poverty. In the present work we decompose political policies along two dimensions—the poor and the rich on one axis and helping and harming on the other. In doing so, we uncover substantial partisan overlap in underlying economic preferences. We find that, counter to some predictions, Democrats do not go out of their way to harm the rich and merely tolerate this as a means to reducing inequality and helping the poor. Moreover, again counter to some predictions, we find that Republicans share Democrats' distaste for providing large gains to the rich. We find broad bipartisan agreement on prioritization of basic economic outcomes (e.g., helping the poor) with differences only emerging once this goal of helping the poor has been met.

Our findings in Study 1 confirmed that people are motivated to some degree by all of helping (versus harming) the poor as well as the rich, reducing (versus increasing) inequality, and increasing (versus reducing) economic efficiency. Of these motives, people were most strongly motivated to help the poor—which was the top priority for Democrats and Republicans alike. However, partisans differ in what they believe should be the focus once the goal of helping the poor is met. In particular, Republicans are strongly averse to harming the rich and thus prefer “rising tide” policies that help both the rich and poor without increasing inequality. Republicans' aversion to harming the rich serves as a constraint on their willingness to reduce inequality: when reducing inequality involves harming the rich, Republicans are less willing to offer their support than Democrats.

Indeed, Democrats strongly prioritize reducing (vs. increasing) inequality and are less opposed to harming the rich in service of that end (particularly when harm to the rich is not extreme). Still, Democrats do not appear to simply harbor strong anti-rich sentiment *per se*: when we explored Democrats' support for helping versus harming the rich across the typology (i.e., holding other factors like support for inequality and efficiency constant), Democrats showed no overall preference for harming over helping the rich (Study 1) or even a slight preference for helping versus hurting the rich (Study 2). When we conducted additional analyses incorporating the *degree* of help and hurt, we discovered that Democrats were averse to helping the rich but *also* to harming them. Taken together, these findings demonstrate that Democrats do not go out of their way to harm the rich, but rather are simply more tolerant of harming the rich as a by-product of achieving their prioritized goals of helping the poor and reducing inequality.

In Study 1 we focused primarily on exploring voting behavior within the typology itself and how these patterns differ along partisan lines. Therefore, in Study 2 we explored the emotional antecedents, and real-world consequences, of partisans' preferences within the typology. The motive to reduce inequality by harming the rich (but keeping the material circumstance of the poor constant) is associated with antipathy towards the rich among both Democrats and Republicans. However, for Democrats, this motive is also associated with compassion for the poor—suggesting that Democrats may believe that alleviating the relative deprivation of the poor is inherently beneficial for them, even if their absolute deprivation is unchanged (33, 34). Lastly, Study 2 showed that people's preferences in the typology reflect within-political party differences (i.e., differences between supporters of Bernie Sanders and Joe Biden), finding that Sanders (relative to Biden) supporters demonstrate greater willingness to harm the rich and concern for the poor as well as less concern about economic efficiency.

Theoretically, our findings offer a new understanding of how people manage tension between desired socioeconomic outcomes. Specifically, we find that people are not simply seeking to maximize the wealth of the poor *or* reduce inequality, but instead weigh different concerns in tandem. Previous research using earlier methods has indicated that, for example, people tend to either wholly make decisions based on maximizing the wealth of the poor (e.g., 20) *or* creating the most egalitarian distribution of resources (e.g., 17). We demonstrate that these considerations are interdependent. For example, we find that both Republicans and, especially Democrats, will forgo some gains to the poor if they come along with steeply rising inequality. If people were purely motivated by increasing the wealth of the poor we would not expect to see such an interaction. People care about not only helping the poor in absolute terms, but also about improving their relative standing.

Practically, this work provides numerous insights into how policy makers might approach framing policies surrounding inequality, poverty, and wealth. Our results suggest that, for example, policies might be more appealing to Democrats to the extent that they focus on helping the poor rather than reducing inequality, and that emphasizing how a policy levies additional taxes upon the rich might actually backfire and reduce policy support. Additionally, our results suggest that more bipartisan support would emerge for policies that emphasize joint gains to both the poor and the rich, so long as the poor are benefitting more than the rich. Our research highlights that there is a great deal more bipartisan agreement than disagreement with respect to economic policy preferences. Framing policies that fall within these areas of agreement (e.g., de-emphasizing harm to the rich, emphasizing help to the poor) is one possible route to effectively fostering bipartisan agreement.

Limitations and Future Directions

This research is not without its limitations. First, in most analyses we utilized model predictions. As indicated in Fig. 3, this is sometimes challenging as we are making inferences beyond the point at which data exist within the typology (e.g., predicting how people would respond to a policy where the rich or poor are being helped greater than 50%). Therefore, one

promising future direction will be to expand the typology to 100% in either direction to get a more rounded picture of policy support at the extremes. Second, while we sampled evenly across political party membership, our sample was not representative. It is crucial to understand how voting within the typology and preferences for outcomes like inequality reduction are reflected with a more comprehensive nationally representative or cross-cultural sample. Thus, moving forward it would be critical to collect these data in such samples. Third, whereas one advantage of our paradigm is that it allows us to ask about a very broad range of policies with carefully controlled effects on the rich and the poor, one downside is that, as a consequence, participants were making decisions about hypothetical policies. Because the participants did not have to directly incur a cost or personally be affected by support for one policy versus another, it is possible that the preferences expressed in our typology would not perfectly match preferences revealed in the real world.

Finally, we might expect that voting behavior, priorities, and how people manage tensions differ depending on the social or economic categories on the axes. While we chose the top and bottom 20% of Americans, we could redefine the axes of the typology in any number of ways. For one, it may be valuable to re-examine our conclusions when we redefine the rich and the poor: for example, it may be that individuals exhibit somewhat different patterns if the rich are re-defined as the top .01%. Beyond this, however, the flexibility of our typology allows us to generalize not only to different operationalizations of rich and poor but also to different contexts (i.e., beyond class divides). For example, we could utilize race- or gender-based axes to examine how individuals manage trade-offs relevant to those contexts—how much inequality between racial or gender groups are people willing to tolerate to improve the material position of the disenfranchised group? One could even redefine help or harm to reflect (for example) social standing or physical harm as opposed to simply material or economic harm. By leveraging our novel and flexible typology structure, and assessing the large array of policies or situations it allows for, one can generate valuable insights about how people might make socioeconomic decisions in the real world whenever multiple motivations are simultaneously at play and in conflict.

Conclusion

Inequality reduction and poverty alleviation are at the center of much social and political debate. Here, we develop and test a typology that delineates the structure of social policy into its key dimensions—the outcome on the poor and the rich. We find, broadly, that policy support derives from a mixture of a desire to reduce inequality, to help the rich, to grow the overall pool of economic resources, and, most strongly, to help the poor. Across partisan lines, we find that Democrats demonstrate a stronger desire to reduce inequality while Republicans demonstrate a stronger desire to avoid harming the rich. Both Democrats' and Republicans' desire to harm (versus help) the rich while holding help to the poor constant, thereby reducing inequality, is related to anti-wealth sentiment. However only Democrats' demonstrate the belief that reducing inequality will inherently benefit the poor. These findings shed light on the psychology of socioeconomic decision making and provide insights for shaping political policies aimed at garnering broad bipartisan agreement.

Materials and Methods

All studies were approved under [BLINDED INSTITUTION'S] Institutional Review Board (IRB), under IRB number STU00211051. All participants provided informed consent. Additionally, all data, materials, and code for Study 1, Study 2, and the supplemental replication of Study 1 can be found online at https://osf.io/rt7uq/?view_only=8d2ae3bc4cc340a9bd525113657a59d4

Study 1

Participants. We recruited a sample of 1,510 American adults from Amazon's Mechanical Turk. Of this initial sample, 53 participants failed an attention check and were thus excluded from analysis leaving a final sample of 1,457 participants. Importantly, the sample was evenly split between Democrats ($n = 728$) and Republicans ($n = 725$). We conducted a priori power simulations (35, 36) which determined that a sample of 1,460 participants offered sufficient power (Power = .99, $\alpha = .05$) to detect standardized regression effects in a binomial crossed-factors multilevel model as small as $\beta = .05$.

Procedure. We presented each participant with a selection of 30 hypothetical political policies randomly drawn from a larger set of 440 policies. Each policy was described as having some effect on the wealth of the country's rich and poor, defined as the top and bottom 20% of Americans. We defined wealth as "the total value of all of a person's financial assets (e.g., investments, savings, homes, etc)." In one such sample policy, participants were told to "imagine a new economic policy was proposed that was projected to (a) result in a 25% increase in the wealth of the country's poor and (b) result in a 5% increase in the wealth of the country's rich. Participants then indicated whether they would vote for this policy (0 = No, 1 = Yes), as well as the degree to which they support this policy on a 7-point Likert scale (1 = "Would Strongly Oppose" to 7 = "Would Strongly Support"). While the question framing remained constant, the percentage change for the rich and poor varied randomly between "no change" and 50% increase or decrease with each policy in 5% increments.

Due to the fact that participants only evaluated 30 (out of 440) randomly selected policies, the resulting data was by design Massively Missing Completely at Random (MMCAR; (32)). In this study the data were 93.1% missing, and the missingness was, by definition, completely at random. The advantage of the MMCAR planned missingness design is that we were able to get a sufficient number of votes per hypothetical policy ($M_{votes} = 102.74$, $Min_{votes} = 95$, $Max_{votes} = 109$) while simultaneously not overwhelming participants with 440 repetitions of the same question. Designing our procedure such that the data is missing completely at random ensures that any missing data techniques lead to unbiased estimates (37).

Following answering the policy questions, participants reported their support for economic inequality ($M = 2.88$, $SD = 1.56$; 6) and social dominance orientation ($M = 2.68$, $SD = 1.35$; 7) on scales ranging from 1 (Strongly [Disagree/Oppose]) to 7 (Strongly [Agree/Favor]). Next, participants reported their gender (55.8% Male, 43.4% Female, 0.7% Other/Prefer not to say), age ($M = 38.82$, $SD = 12.35$), ethnicity (77.4% White, 13.1% Black or African American, 0.8% American Indian or Alaska Native, 6.2% Asian, 0.3% Native Hawaiian or Pacific Islander, 2.2% Other), and household income ($Median = \$50,000$ - $\$59,999$, $SD = 3.48$). Participants then reported their political ideology across four questions. First, participants reported how Liberal (1 = "Very Liberal") or Conservative (7 = "Very Conservative") they are (a) "when it comes to politics ($M = 3.85$, $SD = 1.92$)," (b) "on social issues ($M = 3.70$, $SD = 1.97$)," and (c) "on economic issues ($M = 4.00$, $SD = 1.94$)." Lastly, participants chose "whether [they] most closely identify as a" Democrat ($n = 533$), Independent ($n = 310$), Republican ($n = 576$), or Other ($n = 35$). If a participant chose Independent or Other, we asked the follow-up forced choice question "if you had to choose one, today, which do you more closely identify with," Democrat ($n = 195$) or Republican ($n = 149$). Participants who chose "independent" or "other" initially were included in the category they chose in the second question; prior work demonstrates that those who lean Democrat or Republican, respectively, are psychologically very similar to those who categorize themselves as such (40). See Table S9 for Democrat and Republican Scale Means of key variables.

In order to explore how policy outcomes predicted policy support we ran a series of standardized binomial crossed-factors multilevel regressions. In Model 1 we explored how categorical change (-1 = Decrease, 0 = No Change, and +1 = Increase) in each of the wealth of the poor and rich, inequality, and economic efficiency interacted with political ideology to predict

policy support. In Model 1a we entered only the two grouping variables (participant and hypothetical policy) as well as the four main predictors. In Model 1b we added contrast coded political ideology (-1 = Democrat and 1 = Republican) as well as four interaction terms, one for each of the four main predictors with political ideology. In Models 1c and 1d we recoded political ideology (Model 1c: 0 = Democrat, 1 = Republican; Model 1d: 1 = Democrat, 0 = Republican) in order to get the simple slopes for each main effect among Democrats (Model 1c) and Republicans (Model 1d). In Model 2 we explored how the absolute degree of increase or decrease (0 = No Change to 50 = 50% Change, in increments of 5) to the rich and poor as well as political ideology interacted to predict policy support. In Model 2a we again entered only the two grouping variables (participant and hypothetical policy) as well as the four main predictors. In Model 2b we included contrast coded political ideology (-1 = Democrat and 1 = Republican) as well as all two- and three-way interactions. In Models 2c and 2d we again recoded political ideology (Model 2c: 0 = Democrat, 1 = Republican; Model 2d: 1 = Democrat, 0 = Republican) in order to get the simple slopes for each main effect and Level 2 interaction among Democrats (Model 2c) and Republicans (Model 2d). Additionally, in both Models 1 and 2 we computed variance inflation factors (VIFs) to assess for multicollinearity; all VIFs were below 3, suggesting no issues with multicollinearity (See Tables S10 and S11).

Study 2

Participants. We recruited a sample of 503 American adults from Amazon's Mechanical Turk. Of this initial sample, 29 participants failed an attention check and were thus excluded from analysis leaving a final sample of 474 participants. Importantly, the sample was relatively evenly split between Democrats ($n = 250$) and Republicans ($n = 224$). We conducted a priori power simulations (35, 36) which determined that a sample of 500 participants offered sufficient power (Power = .99, $\alpha = .05$) to detect standardized regression effects in a binomial crossed-factors multilevel model as small as $\beta = .05$.

Procedure. The procedure and policy matrix were nearly identical to Study 1, with three key differences. First, we did not utilize a MMCAR design—participants evaluated each of the presented hypothetical policies. In order to facilitate this complete design, the two additional differences were a restriction of range in the typology. Second, we restricted the policies to a set of 119 that range from “no change” to plus or minus 35%, instead of 50%. Lastly, we removed policies that hurt the poor. We made this final change because from what we observed in Study 1, there was very little support for such policies.

Following this, participants reported whether or not (1 = Yes, 0 = No) they “are a registered Democrat who is able to vote in the Democratic Presidential primary race” ($n_{yes} = 258$). Participants were then asked “who they vote[d] for, or who they [planned] to vote for, in the 2020 Democratic Presidential primary race” out of Joe Biden, Bernie Sanders, Elizabeth Warren, Michael Bloomberg, Tulsi Gabbard, Amy Klobuchar, Pete Buttigieg, or Other. If participants chose “Other” they were presented with the full list of 22 Democratic Candidates. We focused our voter analyses on those who chose Joe Biden ($n = 79$) or Bernie Sanders ($n = 103$).

Next, participants reported their emotions towards the rich and poor. Participants were asked: “In thinking about this country's [poor/rich] (i.e., the people who are in the [bottom/top] 20% of total wealth), how well does each word below describe your feelings towards them?” Specifically, participants reported their feelings of anger, disgust, compassion, and envy using a 5-point scale ranging from 1 (“Does not describe my feelings”) to 5 (“Clearly describes my feelings”). As described in the main text, we composited anger and disgust for each of the rich ($M = 2.41$, $SD = 1.35$) and poor ($M = 1.71$, $SD = 1.10$) into single scores indicating negative emotions. We utilized each of compassion to the rich ($M = 2.06$, $SD = 1.28$) and poor ($M = 3.72$, $SD = 1.16$) as our measure of positive emotions towards the poor. In our main analyses we computed a single emotion score towards each of the poor ($M = 3.72$, $SD = 1.53$) and rich ($M =$

2.06, $SD = 1.49$) in which we subtracted negative emotions from positive emotions towards the poor and rich.

Finally, participants reported social dominance orientation ($M = 2.89$, $SD = 1.31$; 7). Next, participants reported their gender (55.3% Male, 44.5% Female, 0.2% Other/Prefer not to say), age ($M = 37.55$, $SD = 12.38$), ethnicity (81.6% White, 9.3% Black or African American, 0.2% American Indian or Alaska Native, 6.3% Asian, 2.5% Other), and household income (*Median* = \$50,000-\$59,999, $SD = 3.61$). Participants then reported their political ideology across four questions. First, participants reported how Liberal (1 = "Very Liberal") or Conservative (7 = "Very Conservative") they are (a) "when it comes to politics ($M = 4.08$, $SD = 1.92$)," (b) "on social issues ($M = 3.91$, $SD = 1.99$)," and (c) "on economic issues ($M = 4.20$, $SD = 1.90$)." Lastly, participants chose "whether [they] most closely identify as a" Democrat ($n = 195$), Independent ($n = 87$), Republican ($n = 185$), or Other ($n = 6$). If a participant chose Independent or Other, we asked the follow-up forced choice question "if you had to choose one, today, which do you more closely identify with," Democrat ($n = 54$) or Republican ($n = 39$). When analyzing data across partisan lines we categorize those who reported being Democrat or Republican on either question as falling into the respective categories. See Table S12 for Democrat and Republican Scale Means of key variables.

As described in the Results section, we replicated Models 1 and 2 from Study 1. Identical to Study 1, in both Models 1 and 2 we computed variance inflation factors (VIFs) to assess for multicollinearity; all VIFs were below 3.5, suggesting no issues with multicollinearity (See Tables S13 and S14).

Acknowledgments


Funding for this study was provided by the Kellogg School of Management to Nour Kteily. The funder had no role in conceptualization, design, analysis, decision to publish, or preparation of any part of this manuscript.

References

1. R. Wilkinson, K. Pickett, *The spirit level: Why greater equality makes societies stronger* (Bloomsbury Press, 2010).
2. K. Payne, *The broken ladder: how inequality affects the way we think, live, and die* (Viking, 2017).
3. S. Matthews, C. Saraiva, Fed finds race, gender disparities cut U.S. GDP by \$2.6 Trillion. *Bloom. Econ.* (2021).
4. D. Leonhardt, The rich really do pay lower taxes than you. *New York Times* (2019).
5. C. Price, K. Edwards, Trends in Income From 1975 to 2018. *Trends Income From 1975 to 2018* (2020) <https://doi.org/10.7249/wra516-1>.
6. T. Piketty, *Capital in the twenty-first century* (The Belknap Press of Harvard University Press, 2014).
7. A. Giridharadas, *Winners take all: The elite charade of changing the world* (Alfred A. Knopf, 2018).
8. J. E. Stiglitz, *The price of inequality: How today's divided society endangers our future*, Norton & C (2012).
9. R. G. Wilkinson, K. E. Pickett, The enemy between us: The psychological and social costs of inequality. *Eur. J. Soc. Psychol.* **47**, 11–24 (2017).
10. S. Pinker, *Enlightenment now: The case for reason, science, humanism, and progress* (Penguin Books, 2018).
11. H. G. Frankfurt, *On Inequality* (Princeton University Press, 2015).
12. N. G. Mankiw, Defending the one percent. *J. Econ. Perspect.* **27**, 21–34 (2013).
13. R. Agarwal, R. M. Holmes, Let's Not Focus on Income Inequality. *Acad. Manag. Rev.* **44**, 450–460 (2019).
14. A. DiMauro, L. Robinson, The COVID vaccine should be distributed on a first come, first served basis to prevent wasting millions of doses. *Bus. Insid.* (2021).
15. D. Meyer, How to stop vaccine nationalism from prolonging the pandemic. *Fortune* (2021).
16. D. Engelmann, M. Strobel, Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *Am. Econ. Rev.* **96**, 1906–1911 (2006).
17. W. J. Brady, *et al.*, Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 7313–7318 (2017).
18. E. J. Finkel, *et al.*, Political sectarianism in America. *Science (80-.)*. **370**, 533–536 (2020).
19. E. Fehr, K. M. Schmidt, A theory of fairness, competition, and cooperation. *Q. J. Econ.*, 817–868 (1999).
20. B. E. Fehr, M. Naef, K. M. Schmidt, American Economic Association Inequality Aversion , Efficiency , and Maximin Preferences in Simple Distribution Experiments : Comment Author (s) : Ernst Fehr , Michael Naef and Klaus M . Schmidt Source : The American Economic Review , Vol . 96 , No . 5 (. **96**, 1912–1917 (2006).
21. G. E. Bolton, A. Ockenfels, Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: Comment. *Am. Econ. Rev.* **96**, 1906–1911 (2006).
22. T. Kameda, *et al.*, Rawlsian maximin rule operates as a common cognitive anchor in distributive justice and risky decisions. *Proc. Natl. Acad. Sci. U. S. A.* **113**, 11817–11822 (2016).
23. J. Haidt, J. Graham, C. Joseph, Above and below left-right: Ideological narratives and moral foundations. *Psychol. Inq.* **20**, 110–119 (2009).
24. J. T. Jost, E. P. Thompson, Group-based dominance and opposition to equality as independent predictors of self-esteem, ethnocentrism, and social policy attitudes among african americans and european americans. *J. Exp. Soc. Psychol.* **36**, 209–232 (2000).
25. B. J. Lucas, N. S. Kteily, (Anti-)egalitarianism differentially predicts empathy for members of advantaged versus disadvantaged groups. *J. Pers. Soc. Psychol.* **114**, 665–692 (2018).
26. H. Waldfoegel, J. Sheehy-Skeffington, O. Hauser, A. K. Ho, N. Kteily, Ideology selectively shapes attention to inequality. *Proc. Natl. Acad. Sci.*
27. N. S. Kteily, M. D. Rocklage, K. McClanahan, A. K. Ho, Political ideology shapes the

- amplification of the accomplishments of disadvantaged vs. Advantaged group members. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 1559–1568 (2019).
28. G. Mitchell, P. E. Tetlock, B. A. Mellers, L. D. Ordóñez, Judgments of social justice: Compromises between equality and efficiency. *J. Pers. Soc. Psychol.* **65**, 629–639 (1993).
 29. J. Martin, P. Brickman, A. Murray, Moral outrage and pragmatism: Explanations for collective action. *J. Exp. Soc. Psychol.* **20**, 484–496 (1984).
 30. C. J. Wakslak, J. T. Jost, T. R. Tyler, E. S. Chen, Moral outrage mediates the dampening effect of system justification on support for redistributive social policies. *Psychol. Sci.* **18**, 267–274 (2007).
 31. D. Sznycer, *et al.*, Support for redistribution is shaped by compassion, envy, and self-interest, but not a taste for fairness. *Proc. Natl. Acad. Sci.* **114**, 8420–8425 (2017).
 32. W. Revelle, *et al.*, “Web and phone based data collection using planned missingness designs” in *Sage Handbook of Online Research Methods*, 2nd Ed., (Sage Publications, 2016), pp. 578–595.
 33. I. Walker, T. F. Pettigrew, Relative deprivation theory: An overview and conceptual critique. *Br. J. Soc. Psychol.* **23**, 301–310 (1984).
 34. M. Bernstein, F. Crosby, An empirical examination of relative deprivation theory. *J. Exp. Soc. Psychol.* **16**, 442–456 (1980).
 35. P. Green, C. J. MacLeod, simr: an R package for power analysis of generalized linear mixed models by simulation. *Methods Ecol. Evol.* **7**, 493–498 (2016).
 36. T. D. Jorgenson, S. Pornprasertmanit, A. M. Schoemann, Y. Rosseel, semTools: Useful tools for structural equation modeling (2020).
 37. M. Jamshidian, S. Jalal, C. Jansen, MissMech : An R Package for Testing. *J. Stat. Softw.* **56**, 0–31 (2014).
 38. D. Wiwad, *et al.*, The support for economic inequality scale: Development and adjudication. *PLoS One* **14**, 1–29 (2019).
 39. A. K. Ho, *et al.*, The Nature of Social Dominance Orientation: Theorizing and Measuring Preferences for Intergroup Inequality Using the New SDO7 Scale. *J. Pers. Soc. Psychol.* **109**, 1003–1028 (2015).
 40. Pew Research Center, “Appendix A: Why partisans include ‘leaners’” (2015).

Figs. and Tables



Figs. 1a-e. Theoretical typology and heatmaps depicting policy support. (a) Theoretical map of the complete typology, (b) Overall policy support, (c) Policy support among Democrats, (d) Policy support among Republicans, (e) Difference in policy support (Democrat – Republican).

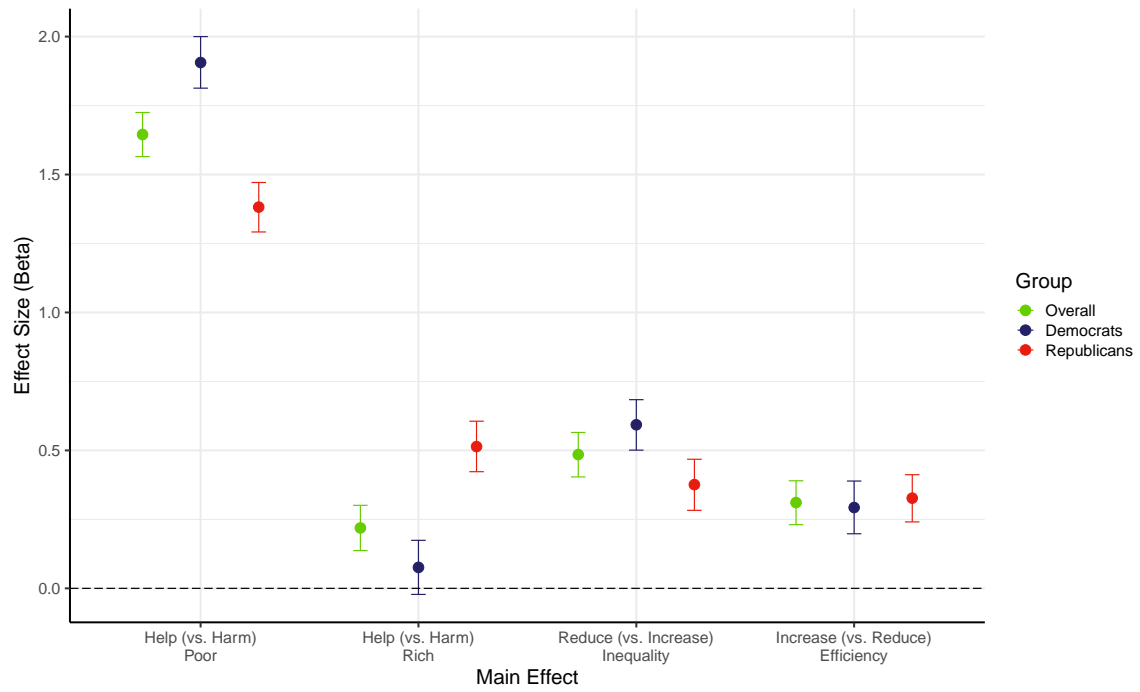


Fig. 2. Beta Weights from Model 1 in Study 1 with 95% Confidence Intervals. All coefficients have been rescored to be positive for ease of comparison. Within-model statistical comparisons between the Overall, Democrat, and Republican coefficients can be found in Table S1.




Fig. 3. Model predictions of policy support among Democrats and Republicans when the poor are being helped but inequality is rising. The blue dot on each line indicates the last point of observed data within our typology; to the right of each blue dot the model is extrapolating beyond existing data (i.e., there are no policies within our typology that help the poor by 10% but also increase inequality by 45% or more).

Table 1. Summary of the multilevel regression analysis for variables predicting binary hypothetical policy voting in Study 1 ($n_{votes} = 43,710$, $n_{participant} = 1,457$, $n_{policies} = 440$). * = $p < .05$, ** = $p < .01$, *** = $p < .001$.

	Model 1a: No interactions		Model 1b: Full Model		Model 1c: Democrats		Model 1d: Republicans	
	β	SE	β	SE	β	SE	β	SE
Intercept (γ_{00})	-0.489	0.048	-0.502	0.048	-0.534 ***	0.065	-0.471 ***	0.064
Fixed Effects								
<u>Level 1 Variables</u>								
Party Identification (γ_{10})	--	--	0.032	0.043	0.063	0.085	-0.064	0.086
<u>Level 2 Variables</u>								
Change to the Poor (γ_{01})	1.603 ***	0.039	1.643 ***	0.041	1.906 ***	0.047	1.382 ***	0.046
Change to the Rich (γ_{02})	0.261 ***	0.040	0.219 ***	0.041	-0.076	0.050	0.514 ***	0.046
Change to Inequality (γ_{03})	-0.493 ***	0.039	-0.484 ***	0.041	-0.593 ***	0.047	-0.376 ***	0.047
Net Loss/Gain (γ_{04})	0.296 ***	0.038	0.310 ***	0.040	0.293 ***	0.049	0.327 ***	0.044
<u>Cross-Level Interactions</u>								
Party*Change to the Poor (γ_{11})	--	--	-0.263 ***	0.022	-0.525 ***	0.045	0.527 ***	0.046
Party*Change to the Rich (γ_{12})	--	--	0.295 ***	0.024	0.590 ***	0.049	-0.590 ***	0.050
Party*Change to Ineq. (γ_{13})	--	--	0.109 ***	0.023	0.218 ***	0.046	-0.218 ***	0.047
Party*Net Loss/Gain (γ_{14})	--	--	0.017	0.022	0.033	0.045	-0.033	0.045
Random Effects								
Intercept – Participant (τ_{00})		2.290		2.284		2.284		2.290
Intercept – Policy (τ_{00})		0.200		0.218		0.218		0.218

Note. Party Identification in Model 1b is contrast coded where Democrat is coded as -1 and Republican is coded as 1. Party Identification in Model 1c is coded with Democrats as 0 and Republicans as 1 in order to get simple slopes for Democrats. The reverse is true in Model 1d.