The Policy-People Gap:

Decision Makers Choose Policies That Would Select Different Applicants Than They Select When Making Individual Decisions

SUPPLEMENTARY MATERIALS

The supplement is structured as follows:

- In "SM 1. Main Studies: Expanded Methods and Results" we describe additional measures and analyses for the studies reported in the manuscript
- In "SM 2. Supplemental Studies" we describe the methods and results of fifteen supplemental studies.
- In "SM 3. Results of All Non-Choice Measures across Studies" we report results comparing policy and individual decisions for all of the non-choice measures we collected. When describing the measures, we provide a label to facilitate matching the measure with its results.
- In "SM 4. Other Analyses" we report (1) analyses of whether various individual differences (e.g., gender, political attitudes) moderate the policy-people gap and (2) a summary of alternative mechanisms we tested for the policy-people gap.
- In "SM 5. Supplementary References" we provide the references cited here.

All of our data, syntax, and survey materials are available here: https://bit.ly/3tJB1q7

TABLE OF CONTENTS

SM 1. MAIN STUDIES: EXPANDED METHODS AND RESULTS

SM 1.0. INTRODUCTION AND GENERAL ANALYSES

Introduction to additional measures

We use short labels (label: *Sample Label*) to refer to measures we collected across the studies. We provide these labels to facilitate linking the descriptions of the measures with the results presented in SM 3. In SM4 we report a construal-level analysis of the qualitative measures collected.

SM 1.1. STUDY 1 (COLLEGE ADMISSIONS)

Additional measures

In addition to the measures reported in the main manuscript, participants answered a series of exploratory questions about their choices. Specifically, for five different factors, participants reported for which of the two decisions they made (individual or policy decision) they more focused on that factor (-5 = *Much more focused on this for the policy decision*, 0 = *Equally focused*, 5 = *Much more focused on this for the specific applicant decision*). The five factors were (label: *Focus: Objective Information*) "the objective achievements associated with each of the applicants," (label: *Focus: Values*) "your values and the values of your institution," (label: *Focus: Applicant Adversity*) "the adversity that one or more applicant(s) faced," (label: *Focus: Justification*) "how easy it would be to justify your choice," and (label: *Focus: Identity*) "how your choice reflected your identity." For each of these items we conducted a one-sample t-test against the midpoint of the scale (0). Results for these items are reported in SM 3.

Additional analyses

In addition to the analyses reported in the main manuscript, we analyzed whether the order of the decisions (i.e., whether the individual or policy decision came first) influenced choice. For this analysis, we necessarily examine all choices (Time 1 and Time 2). We conducted a logistic regression on choice (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income option) using Decision Type (1 = Individual, 0 = Policy), Order (1 = Individual first, Policy second; 0 = Policy first, Individual second), and the Decision Type x Order interaction as independent variables, and we clustered standard errors at the participant level. We found a significant coefficient of Decision Type ($b = 1.085$, $p < .001$), such that participants were more likely to select the higher-scoring, higher-income option in the Individual condition when the Policy decision came first. We did not find a significant coefficient of Order ($p = .594$), nor

of the Decision Type x Order interaction (p = .686), suggesting that the preference for the higher-scoring, higher-income option in the Individual condition is not moderated by whether the Individual or Policy decision came first.

We also conducted a logistic regression of Fair Choice (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income) using the same independent variables and specifications as in the Choice regression. We find the same pattern of results as with Choice: a significant coefficient of Decision Type (b = 0.855, p < .001), no significant coefficient of Order (p = .801), and no significant Decision Type x Order interaction (p = .974).

We conducted an OLS regression of Likelihood (higher values indicate a greater likelihood of selecting the higher-scoring, higher-income option) using the same independent variables and specifications as in the choice regression. We find the same pattern of results as with Choice and Fair Choice: a significant coefficient of Decision Type (b = 1.103, p < .001), no significant coefficient of Order (p = .632), and no significant Decision Type x Order interaction (p = .708).

SM 1.2. STUDY 2 (WORKPLACE HIRING)

FIGURE S1
Stimuli (Study 2)

	Candidate A	Candidate B
Coding Assignment Score (out of 10)	8.1	8.0
Undergraduate Grade Point Average (out of 4.0)	3.8	3.8
References	2 (strong)	2 (strong)
Candidate Gender	Male	Female

Note. In the Individual condition we labeled the columns "Candidate [A|B]" and in the Policy condition we labeled them "Candidate Type [A|B] (Policy [A|B])". We counterbalanced whether the higher-scoring, male option or the lower-scoring, female option was labeled "A" and appeared on the left.

Additional measures

In addition to the measures reported in the main manuscript, participants answered a series of exploratory mechanism items. First, participants reported their beliefs about the effect of gender on outcomes. All participants answered the same pair of questions on 7-point Likert scales (1 = *Not at all*, 4 = *Moderately*, 7 = *Extremely)*, but the phrasing of the questions varied slightly depending on which type of decision participants made earlier. Participants who made individual decisions reported their beliefs about the specific candidates saw (e.g., they reported how much "these candidates' gender affected" their scores), whereas participants who made policy decisions reported their beliefs about candidates in general (e.g., they reported how much "candidates' gender affects" their scores). Participants answered the following two questions (this is the Policy version of the questions): "To what extent do you think candidates' gender affects their Coding Assignment Scores?" and "To what extent do you think gender causes

differences in candidates' Coding Assignment Scores?". Per our preregistration, we created a composite of these two items (label: *Situational Attribution; r* = .80).

Participants also indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with the three moralization items: "My choice of [candidate/policy] is connected to my beliefs about fundamental right and wrong.", "There are very important ethical aspects to this choice.", and "My choice of [candidate/policy] is tied to my core moral values and beliefs." Per our preregistration, we created composite measure of these items (label: *Moralization composite;* α = .79). Next, participants answered two exploratory questions. They reported how much "differences in skill or effort cause differences in candidates' Coding Assignment Scores" (label: *Skill/Effort*; this question had the same format as the questions about the effect of gender) and provided an open-ended answer for why they made their choice of candidate/policy. Results are reported in SM 3.

Additional analyses

We tested two mediation models (referred to as 1a and 2a in the pre-registration). Model 1a examined whether drawing situational attributions for the candidate's scores (in this case, seeing the scores as a reflection of the candidates' gender) mediated the effect of making an individual versus policy decision on choosing the higher-scoring, male option. Model 2a examined whether moralizing the choice (that is, seeing the choice as being related to one's moral values and beliefs) mediated the effect of making an individual versus policy decision on choosing the higher-scoring, male option. For both of these models we conducted binary mediation analyses, setting the choice of candidate/policy as the dependent variable (1 = higher-scoring, male, 0 = lower-scoring, female), the type of decision as the independent variable, (1 = Individual, 0 = Policy or Rule), and the Situational Attribution composite or Moralization

composite as the mediator. We ran the two models using the PROCESS Macro for SPSS (Model 4, Hayes, 2018), which allows us to estimate the indirect effect using a percentile bootstrap estimation approach with 10,000 samples (Shrout & Bolger, 2002).

Model 1a: Situational attribution. We found that drawing a situational attribution for the candidates' scores did not significantly mediate the relationship between making an individual versus policy decision and the choice of candidate policy (indirect effect = 0.106, *SE* = 0.060; 95% *CI* = [-0.009, 0.230]). That said, participants making an individual decision (relative to a policy decision) were marginally less likely to draw a situational attribution ($a = 0.245$, $p = .072$) and drawing a situational attribution was associated with a significantly lower likelihood of selecting the higher-scoring, male option ($b = -0.423$, $p < .001$). The relationship between type of decision and choice was marginally significant when the mediator was included in the model ($c' = 0.424$, $p = .058$). These results raise the possibility that situational attributions could be affected by making policy versus individual decisions. We encourage future research to investigate the relationship between policy decisions and causal attributions, especially because we find evidence that situational attributions influence people's selection decisions.

Model 2a: Moralization. We found that moralizing the choice did not significantly mediate the relationship between making an individual versus policy decision and the choice of candidate policy (indirect effect = 0.039, *SE* = 0.037; 95% *CI* = [-0.015, 0.127]). Participants making an individual decision (relative to a policy decision) were not significantly more or less likely to moralize the choice ($a = -0.194$, $p = .178$). Moralizing the choice was associated with a marginally-significant lower likelihood of selecting the higher-scoring, male option ($b = -0.203$, $p = .014$). The relationship between type of decision and choice was marginally significant when the mediator was included in the model ($c' = 0.464$, $p = .033$). These results suggest that

decision-makers' *perceptions* of how much individual and policy decisions relate to their values are not significantly different.

SM 1.3. STUDY 3 (MEDIATION BY SHIFTING STANDARDS OF FAIRNESS)

Additional analyses

Socially desirable responding. We measured socially desirable responding (SDR) through two abridged versions of common scales, the Marlowe-Crowne Social Desirability Scale (MCSDS; Crowne & Marlowe, 1964) and the Balanced Inventory of Desirable Responding (BIDR; Paulhus, 1988). The MCSDS scale "represents a single latent construct—need for approval," while the BIDR represents "impression management (IM) and self-deceptive enhancement" (pp. 1-2, Hart et al., 2015). Participants completed the M-C 1 (Strahan & Gerbasi, 1972) and the BIDR-16 scale (Hart, Ritchie, Hepper, & Gebauer, 2015); these abridged scales are recommended among alternatives because of their reliability and correlation with their unabridged versions (Fischer & Fick, 1993; Hart, Ritchie, Hepper, & Gebauer, 2015). We scored the BIDR-16 continuously (as recommended by Stober, Dette, & Musch, 2002): inversely keyed items were reversed, and points associated with each answer were summed across items. We scored the MCSDS scale as standard: summing the number of items where participants gave the socially desirable response, from 0 to 10.

In order to test whether SDR explains the observed divergence, we tested whether there is a significant association between measures of SDR and the choice of who to admit. Social desirability for all three scales did not vary based on making a policy versus individual decision: BIDR-IM (p = .904), BIDR-SDE (p = .207), nor MCSDS (p = .246). A regression (R^2 = .024) predicting choice with factors for condition and all three scales revealed significant coefficients for condition (coeff. = .108, p = .002) and BIDR-IM (coeff. = .007, p = .004), but not for BIDR-SDE (p = .655), nor MCSDS (p = .204). In sum, compared to a regression with just condition as

a factor (R^2 = .012; coeff. = .112, p = .002), including the social desirability scales improved prediction and did not change the effect of condition.

Choice mediation split by order. We conducted the same mediation reported in the manuscript, but split the data based on whether participants answered the mechanism questions before or after making a choice. As noted in the manuscript, when participants answered the mechanism questions before making a choice, we did not find evidence of mediation. We found that participants making the Individual decision, compared to the Policy one, were only marginally more motivated by microjustice concerns (b = 0.08, p = .099), which were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option (b = 0.51, p < .001). Furthermore, participants making the Individual decision, compared to the Policy decision, were not significantly more or less motivated by macrojustice concerns (b = -0.03, p = .866), although these were significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -0.95, p < .001). The confidence intervals for the microjustice and macrojustice indirect effect did not omit zero (microjustice indirect effect = 0.08, 95% *CI* = [-0.02, 0.21]; macrojustice indirect effect = 0.03, 95% *CI* = [-0.30, 0.37]).

When participants answered the mechanism questions after making a choice, we found the expected mediation results. We found that participants making the Individual decision, compared to the Policy one, were significantly more motivated by microjustice concerns (b = 0.16, p = .001), which were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option (b = 0.95, p < .001). Furthermore, participants making the Individual decision, compared to the Policy decision, were also significantly less motivated by macrojustice concerns (b = -0.50, p = .016), which were significantly associated with a lower

likelihood of choosing the higher-scoring, higher-income option (b = -1.44, p < .001). The confidence intervals for both the microjustice and macrojustice indirect effects omitted zero (microjustice indirect effect = 0.36, 95% *CI* = [0.14, 0.66]; macrojustice indirect effect = 0.72, 95% *CI* = [0.13, 1.40]).

SM 1.4. STUDY 4 (AN INTERVENTION TO ALIGN CHOICES)

Additional analyses.

Mediation analyses. Per our preregistration, we tested an additional mediation model in which we replaced the macrojustice and microjustice items with a single overall justice composite (alpha = .73). However, we urge caution in interpreting these results because of the measure's low reliability and because the concepts of microjustice and macrojustice are not strongly negatively correlated.

Choice mediation, overall composite. Control participants making the Individual decision, compared to the Policy one, were significantly less motivated by the microjustice-to-macrojustice composite (b = 0.66, p < .001), which was significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -1.44, p < .001; Control justice indirect effect = 0.95, 95% *CI* = [0.64, 1.29]). For Macrojustice participants, the indirect effect was smaller than for the Control participants, yet still omitted zero (Macro justice indirect effect = 0.39, 95% *CI* = [0.05, 0.75]; index of moderated mediation = -.55, 95% *CI* = [-1.02, -0.11]). As in the other mediations, the index of moderated mediation suggests that the Macro intervention attenuated the gap in choice between individual and policy decisions because it reduced – though did not completely eliminate – the degree to which decision makers' standards of fairness shifted across the decision contexts.

SM 2. SUPPLEMENTAL STUDIES

SM 2.0. INTRODUCTION AND GENERAL ANALYSES

In SM 2, we describe supplemental studies and report their results. All choice and fair choice results are reported in Tables S2 and S3, and detailed analyses are reported in each study's section. Results from all other measures where we compare the Individual and Policy conditions are reported in SM 3. To be consistent and succinct, we primarily relied on t-tests in our analyses and omitted redundant preregistered analyses (i.e., omnibus ANOVAs in Study S3 and ANOVAs in Studies S12 and S13).

TABLE S1
Percent Selected the Higher-Scoring, Advantaged Option over the Lower-Scoring, Disadvantaged Option (All Supplemental Studies)

Study	Context	*N*	Comparison	Individual	Policy	χ^2	*p*
S1[a]	Admissions	2,016	Standard	70.08%	39.00%	196.42	<.001
S2[b]	Admissions	1,009	Standard				
S3	Admissions	603	Policy-specific	73.33%	35.35%	29.69	<.001
			Policy-range		30.48%	38.63	<.001
S4	Admissions	807	Policy-Narrow	64.73%	51.26%	7.57	.006
			Policy-Broad		48.24%	11.24	.001
			Policy-Abstract		40.59%	23.91	<.001
S5	Hiring	859	Policy-Narrow	73.70%	72.79%	27.23	.806
			Policy-Abstract		62.72%	32.19	.005
S6	Admissions	593	Standard	63.13%	36.41%	27.23	<.001
			Rule		35.07%	32.19	<.001
S7	Admissions	283	Policy-Narrow	56.77%	43.75%	4.76	.029
S8	Admissions	238	Policy-Narrow	59.84%	45.13%	5.19	.023
S10	Admissions	595	Adversity Not Salient	73.83%	48.03%	21.01	<.001
			Adversity Salient	46.26%	44.22%	0.13	.725
S11	Admissions	595	Control	67.02%	40.21%	13.79	<.001
			Objective Information	82.47%	69.70%	4.39	.036
			Values	47.62%	29.13%	7.51	.006
S12	Admissions	298	Standard	60.96%	53.29%	1.79	.181
S13	Admissions	287	Standard	68.67%	46.72%	14.18	<.001

Note. [a] indicates average choice across all stimuli
[b] indicates average choice across all (9) decision rounds.

TABLE S2
Percent Indicated the Higher-Scoring, Advantaged Option over the Lower-Scoring, Disadvantaged Option as the Fair Choice (All Supplemental Studies)

Study	Context	*N*		Individual	Policy	χ^2	*p*
S3	Admissions	603	Policy-specific	80.19%	50.00%	40.28	<.001
			Policy-range		40.69%	67.15	<.001
S4	Admissions	807	Policy-Narrow	71.01%	55.28%	10.82	.001
			Policy-Broad		48.74%	20.99	<.001
			Policy-Abstract		34.65%	54.27	<.001
S5	Hiring	859	Policy-Narrow	84.43%	76.33%	5.96	.015
			Policy-Abstract		65.85%	26.62	<.001
S6	Admissions	593	Standard	70.20%	32.07%	55.56	<.001
			Rule		26.07%	79.80	<.001
S7	Admissions	283	Policy-Narrow	63.23%	32.81%	25.94	<.001
S8	Admissions	238	Policy-Narrow	62.99%	39.82%	12.86	<.001
S10	Admissions	595	Adversity Not Salient	75.17%	40.79%	36.46	<.001
			Adversity Salient	59.18%	46.94%	4.43	.035
S11	Admissions	595	Control	69.15%	49.48%	7.46	.006
			Objective Information	69.07%	47.47%	9.39	.002
			Values	64.76%	47.57%	6.24	.012
S13	Admissions	287	Standard	71.33%	37.23%	33.66	<.001

Relation between constructs of macrojustice and microjustice to constructs of moralization/values and objective information.

Across several studies (S3-S6 and S10-S13) we collected measures pertaining to moralization/values and objective information. Specifically we measured the extent to which participants saw their choice as moralized or pertaining to their values, as well as how much their choice relied on objective information. These measures reflect earlier theorizing around what might underlie the policy-people gap. Overall, across these studies we found that participants saw the policy decision as relying more on their values and less on objective information than the individual decision. However, we believe our present theoretical account is consistent with these results, and better accommodates the fairness results (i.e., that participants making both decisions

saw their decisions as being fair, which would suggest that they both were related to values in some way). Our theoretical account is consistent with these items and results because they overlap with the notions of macrojustice and microjustice. The principles of macrojustice that are relevant to the selection contexts we focus on pertain to equality and diversity. To the extent that participants relied on these values (as our theorizing around macrojustice would suggest) and were aware, would they have indicated that their decision was moralized or relied on their values. In the case of microjustice, participants who sought to rely on applicants' qualities to determine what they ought to get (as microjustice suggests) likely attended to what participants would describe as constituting (seemingly) objective information – test scores.

SM 2.1. STUDY S1 (STIMULUS SAMPLING)

In Study S1, we tested whether the gap between decisions about policies and about individuals replicated across a range of stimuli. To test this proposition and identify potential boundaries of the gap, we modified the scores and incomes of the Study 1 applicant profiles. We paired one of five higher-scoring applicant profiles with one of five lower-scoring applicant profiles, in order to test 25 unique combinations of profiles. We preregistered this study on AsPredicted.org (https://aspredicted.org/NMM_BTG).

Methods

Participants. We recruited online participants on Amazon's Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition. Our final sample comprised 2,016 participants (M_{age} = 37.56, *SD* = 11.62; 55.48% female).

Procedure and materials. We replicated the procedure of Study 1 with two key modifications. We randomly assigned participants to one of fifty conditions based on a 2 (Decision Type: Individual, Policy) x 5 (Higher-scoring, higher-income applicant's profile) x 5 (Lower-scoring, lower-income applicant's profile) between-subjects design. See Figure S2 for the applicant profiles we tested. As in the earlier studies, all participants were asked to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. As before, participants choosing between individuals examined a table of information describing the specific applicants, whereas participants choosing between policies examined a table describing the applicants that would be admitted by each policy. The Policy condition indicated that the policy would apply "to these two applicants and other applicants like them in the future". We presented the table of information as in Study 1: the

lower-scoring, lower-income option was always shown on the left and referred to as "A" and the higher-scoring, higher-income option on the right and referred to as "B".

We generated the 25 applicant profiles based on three sources: the official 2016 SAT score report, the official 2015 SAT score data, and estimates for 2018 household income brackets[1]. We created profiles based on the correlation between income and SAT scores (as is evident in the 2016 report) and made the scores and income vary linearly in equal increments in order to facilitate our analysis. To create the profiles we established nine income categories and paired them with their corresponding SAT score (based on the 2016 report). We adjusted the scores in each category to be one standard deviation higher than the category mean and we rounded them to vary by exactly 60 SAT score points (i.e., the average score difference between all categories) from the category immediately above and below it. We matched the SAT scores with their appropriate percentiles (using the 2015 SAT score data[2]) and the incomes with their appropriate percentiles (using the 2018 estimates for household income brackets).

After participants read about each applicant or policy, we asked them to select one of the two applicants or policies (forced choice: A, B) and to indicate which applicant or policy was the fairer one to select (forced choice: A, B). Participants also indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with three statements from Feldman's (1988) Equality of Opportunity items: "If people were treated more equally in this country, we would have many fewer problems," "Our society should do whatever is necessary to make sure that everyone has

[1] We used readily available estimates (https://web.archive.org/web/20190915204121/https://dqydj.com/average-median-top-household-income-percentiles/).

[2] We used the 2015 SAT score percentiles because that was the last year that the College Board released the percentiles that correspond to every score. The percentiles are likely to still be accurate to the 2016 scores within a small margin of error (https://web.archive.org/web/20160629222935/http://blog.prepscholar.com/sat-percentiles-high-precision-2016).

an equal opportunity to succeed," and "One of the big problems in this country is that we don't give everyone an equal chance." At the end of the survey, participants explained in writing why they made their choice of applicant/policy and reported demographic information.

FIGURE S2
Applicant Profiles (Study S1)

Lower-Scoring, Disadvantaged		Top-Scoring, Advantaged	
SAT	Income	SAT	Income
1630	$10,000	1930	$110,000
1690	$30,000	1990	$130,000
1750	$50,000	2050	$150,000
1810	$70,000	2110	$170,000
1870	$90,000	2170	$190,000

Note. Participants saw one applicant profile from the left column paired with one applicant profile from the right column. The applicant profiles were presented in the same format as in Study 1.

Results

Choice. We found a significant effect of Decision Type on choice (χ^2 [1, 2016] = 196.42, $p < .001$), collapsing across applicant profiles. Participants chose the higher-scoring, higher-income option in the Individual condition more (70.08%) than in the Policy condition (39.00%).

We conducted three logistic regressions on Choice (1 = higher-scoring, higher-income option; 0 = lower-scoring, lower-income option), first using just Decision Type (1 = Individual, 0 = Policy) as an independent variable, then including the higher-scoring, higher-income applicant's SAT score (Profile 1) and the lower-scoring, lower-income applicant's SAT score (Profile 2) as independent variables, and with all interactions. We mean centered and divided the SAT scores by 60 to arrive at more interpretable regression coefficients. Hence, the coefficients for the profiles correspond to a 60-point change in SAT scores (i.e., the difference between each profile's scores). We present the regression results in Table S3 and depict results in Figure S3.

We found a main effect of Decision Type across all three regressions, indicating that participants were more likely to favor the higher-scoring, higher-income option when choosing between individuals than policies, even after controlling for the applicant profiles. Participants were more likely to select the higher-scoring applicant in the Individual condition than in the Policy condition in all but one of the 25 combinations of profiles. We generally saw a small or no effect of applicant profile. However, we did find a Profile 1 (Higher-scoring, higher-income applicant) x Decision Type interaction such that we see a larger difference between the decisions when the higher-scoring, higher-income applicant's scores and income were particularly high.[3]

FIGURE S3
Choice of the Higher-Scoring, Higher-Income Option across Stimuli (Study S1)

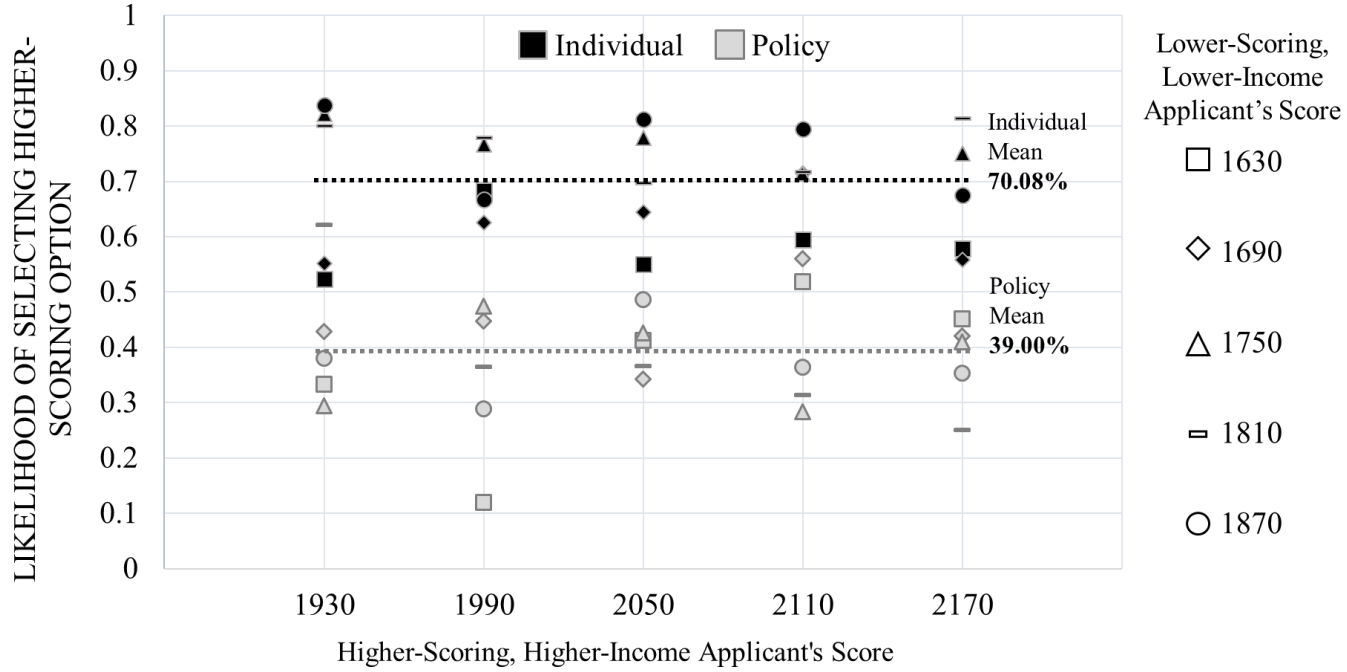

Note. The x-axis indicates the score of the higher-scoring, higher-income applicant. The legend symbols indicate the score of the lower-scoring, lower-income applicant. The black data points correspond to the Individual condition and the gray ones correspond to the Policy condition.

[3] Because the applicants' scores were (by design) perfectly related to their income, it is impossible to say whether the effect of the applicants' profiles is due to participants' perception of the applicants' score or their income, or both.

TABLE S3
Results of Logistic Regression Predicting Choice of Higher-Scoring, Higher-Income Option (Study S1)

	Baseline				Applicant Profile controls				Decision Type x Applicant Profile			
Variable	***b***	***SE***	***z***	***p***	***b***	***SE***	***z***	***p***	***b***	***SE***	***z***	***p***
Decision Type (Individual = 1, Policy = 0)	1.298	0.094	13.76	< .001	1.308	0.095	13.81	<.001	1.328	0.096	13.89	<.001
Profile 1 Score (Higher-scoring)					0.081	0.034	2.39	0.017	-0.057	0.048	-1.20	0.229
Profile 2 Score (Lower-scoring)					-0.012	0.034	-0.36	0.721	-0.006	0.047	-0.12	0.902
Profile 1 x Profile 2					-0.053	0.024	-2.16	0.031	-0.081	0.034	-2.38	0.017
Decision Type x Profile 1									0.288	0.069	4.19	<.001
Decision Type x Profile 2									-0.027	0.068	-0.39	0.693
Decision Type x Profile 1 x Profile 2									0.059	0.049	1.18	0.24
Constant	-0.447	0.065			-0.453	0.065			-0.458	0.065		
Cox & Snell R^2		0.094				0.099				0.108		

Note: Profile 1 Score refers to the higher-scoring, higher-income applicant's SAT score and Profile 2 Score to the lower-scoring, lower-income applicant's SAT score. We mean centered and divided the SAT scores by 60 to arrive at regression coefficients that correspond to the difference in scores between the profiles. This means that the coefficients for the profiles correspond to a 60-point change in SAT scores.

TABLE S4
Results of Logistic Regression Predicting Fair Choice of Higher-Scoring, Higher-Income Option (Study S1)

	Baseline				Applicant Profile controls				Condition x Applicant Profile			
Variable	*b*	*SE*	*z*	*p*	*b*	*SE*	*z*	*p*	*b*	*SE*	*z*	*p*
Condition (Individual = 1, Policy = 0)	1.850	0.099	18.67	< .001	1.852	0.099	18.67	<.001	1.863	0.100	18.69	<.001
Profile 1 Score (Higher-scoring)					0.020	0.036	0.57	0.572	-0.061	0.050	-1.23	0.218
Profile 2 Score (Lower-scoring)					-0.008	0.035	-0.22	0.827	-0.070	0.049	-1.43	0.153
Profile 1 x Profile 2					-0.020	0.025	-0.78	0.435	-0.042	0.036	-1.19	0.235
Condition x Profile 1									0.163	0.071	2.28	0.022
Condition x Profile 2									0.125	0.071	1.76	0.078
Condition x Profile 1 x Profile 2									0.048	0.051	0.94	0.349
Constant	-0.772	0.068			-0.774	0.068			-0.781	0.068		
Cox & Snell R^2	0.175				0.175				0.179			

Note: Profile 1 Score refers to the higher-scoring, higher-income applicant's SAT score and Profile 2 Score to the lower-scoring, lower-income applicant's SAT score. We mean centered and divided the SAT scores by 60 to arrive at regression coefficients that correspond to the difference in scores between the profiles. This means that the coefficients for the profiles correspond to a 60-point change in SAT scores.

Fair choice. We found a significant overall effect of Decision Type on participants' fair choice judgment (χ^2 [1, 2016] = 374.45, $p < .001$), collapsed across applicant profiles. Overall, participants making a decision about individuals were more likely to indicate that the higher-scoring, higher-income applicant (74.61%) was the fair choice than participants making a choice about policies (31.60%). We conducted the same set of logistic regressions on participants' Fair Choice judgment as we did on Choice and we found the same pattern of results (see Table S4).

Discussion

We found that people are more likely to favor the higher-scoring, advantaged applicant over the lower-scoring, disadvantaged applicant across varying levels of scores and income for each applicant. In particular, the gap between policy and individual decisions replicated in 24 out of the 25 unique combinations of alternatives we tested.

SM 2.2. STUDY S2 (REPEATED DECISIONS)

In Study 1 we found that decision makers chose admissions policies that diverged from the specific applicants they chose to admit. However, we found this evidence by comparing the policy decision to a single decision about which applicant to admit. It is possible that over repeated individual decisions, decision makers' choices about which individual applicants to admit might converge with their policy preferences. This could happen because, for example, decision makers realize the cumulative effects of their decisions (Fishbach & Converse, 2010) or they wish to seek more variety as they make more decisions (Read & Loewenstein, 1995). To test this proposition, we had participants in Study S2 make multiple decisions: an initial individual decision followed by eight additional individual decisions or an initial policy decision followed by eight additional decisions to fine-tune their policy. This allowed us to examine whether the composition of decision makers' choices across multiple decisions diverged. We preregistered this study on AsPredicted.org (https://aspredicted.org/HDZ_QTX).

Methods

Participants. We recruited online participants on Amazon's Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition. Our final sample comprised 1,009 participants (M_{age} = 36.89, *SD* = 12.40; 54.71% female).

Procedure and materials. We randomly assigned participants to one of two conditions (Decision Type: Individual, Policy). Participants made a series of nine choices in each condition. When making their initial individual or policy decision (Round 1), participants did not know that they would be making additional decisions (Rounds 2-9). Hence, the initial individual and policy decision had the same experimental format as Study 1. After making their initial decision, participants were instructed that they would now be making "a series of additional decisions" in

which they had to choose between two applicants. Participants in the Policy condition were further instructed that they would make "a series of additional decisions that allow you to fine-tune your policy. Specifically, you will make 8 additional decisions about who your policy will admit." We chose to add this "fine-tune" language because we wanted participants to feel free to select any applicant in the ensuing decisions, rather than feel pressured to be consistent with the policy they selected (e.g., by exclusively choosing lower-scoring applicants in ensuing decisions because the policy they selected favored the lower-scoring type of applicant).

We developed new applicant profiles for the nine decision rounds. Specifically, we used 9 profiles that we developed for Study S1. From the applicant profiles in Study S1, we combined the three highest-scoring profiles from the lower-scoring applicant profiles with the three highest-scoring profiles from the higher-scoring applicant profiles. The rate at which either option was chosen was closest to 50% for these combinations, which is desirable to avoid any ceiling or floor effects.

The initial individual or policy decision (Round 1) always paired the average lower-scoring applicant profile (1810 SAT score, $70,000 household income) with the average higher-scoring applicant profile (2110 SAT score, $170,000 household income). Hence, the initial decision captured the average profile combination that participants would make a decision about. For the ensuing eight decisions (Rounds 2-9), participants were presented with the profile combinations in random order. Thus, the specific applicant profile that some participants saw in Round 2, other participants saw in Round 3, others in Round 4, etc. For each decision that participants made, we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left. For each decision, participants indicated which applicant they wanted to admit (in the Individual condition) or

which applicant they wanted their policy to admit (in the Policy condition). After making all nine decisions, participants provided demographic information (age, gender, and political attitudes).

Results

Choice for Round 1. Participants making their initial choice (Round 1) about individuals were more likely to select the higher-scoring, higher-income option (71.15%) than participants making their initial choice about policies (49.50%; χ^2 [1, 1009] = 49.38, $p < .001$), consistent with Hypothesis 1.

Choices for Rounds 1-9.[4] Participants who made an initial individual decision selected higher-scoring, higher-income applicants across their nine decisions 73.73% of the time, on average, whereas participants who made an initial policy decision selected policies that would favor those applicants only 57.23% of the time (t[1007] = 6.56, $p < .001$, $d = 0.41$).

We conducted an OLS regression predicting participants' choice (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income option), with Decision Type (1 = Individual, 0 = Policy), Decision Round (1-9), and the Decision Type x Round interaction as factors, clustering standard errors at the participant level. We depict the results of the regression in Table S5. As depicted in Figure S4, participants who first made an individual decision were more likely to select the higher-scoring option, both overall and across all decision rounds, than participants who initially made a policy decision.

[4] Per our preregistration, we also conducted two additional analyses. First, we used a chi-square test to examine whether the proportion of participants who chose the higher-scoring, higher-income applicant across the 9 rounds varied based on Decision Type (χ^2 [1, 1009] = 49.55, $p < .001$). However, the t-test we report here may be more appropriate because our dependent variable is quasi-continuous.

TABLE S5
Results of OLS Regression Predicting Choice of Higher-Scoring, Higher-Income Option (Study S2)

Variable	*b*	*SE*	*t*	*p*
Decision Type (Individual = 1)	0.216	0.030	7.19	<.001
Decision Round				
Round 2	0.087	0.019	4.68	<.001
Round 3	0.085	0.018	4.70	<.001
Round 4	0.095	0.018	5.19	<.001
Round 5	0.095	0.019	5.07	<.001
Round 6	0.080	0.018	4.34	<.001
Round 7	0.082	0.018	4.65	<.001
Round 8	0.083	0.019	4.41	<.001
Round 9	0.087	0.018	4.91	<.001
Decision Type x Decision Round				
Individual x Round 2	-0.046	0.024	-1.94	0.053
Individual x Round 3	-0.042	0.024	-1.74	0.083
Individual x Round 4	-0.091	0.024	-3.79	<.001
Individual x Round 5	-0.070	0.024	-2.90	0.004
Individual x Round 6	-0.046	0.023	-1.99	0.047
Individual x Round 7	-0.050	0.024	-2.07	0.039
Individual x Round 8	-0.058	0.024	-2.38	0.017
Individual x Round 9	-0.060	0.023	-2.57	0.010
Constant	0.495	0.022		
Clusters	1,009			
R^2	0.032			
$F(17,1008)$	5.87, $p < .001$			

Note: Standard errors are clustered at the participant level.

FIGURE S4
Choices across Rounds (Study S2)

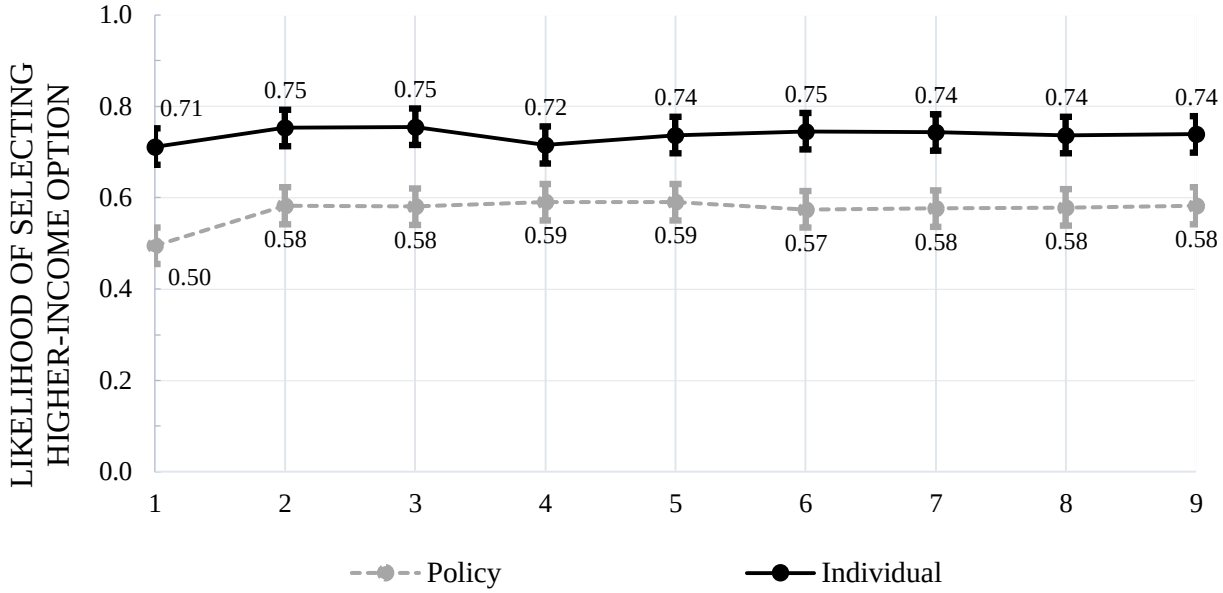

Note. X-axis indicates the decision round. Error bars show 95% confidence intervals (CIs).

Discussion

Even when making numerous decisions, participants were more likely to admit higher-scoring, advantaged applicants over lower-scoring, disadvantaged applicants when choosing between individuals than when choosing between policies, consistent with Hypothesis 1. The number of higher-scoring, advantaged applicants that participants admitted over multiple, individual decisions (73.73%) did not converge with the number of participants who favored an initial policy that would admit those applicants (49.50%), nor with the overall policy that participants defined across multiple decisions (57.23%). This result, therefore, documents the existence of a gap between policy and individual decisions even when comparing policies to the overall selection of applicants across multiple individual decisions.

SM 2.3. STUDY S3 (SPECIFIC VS RANGE STIMULI)

In Study S3, we tested whether the divergence between individual and policy decisions is robust to (a) alternative policy descriptions and (b) alternative information about the applicants. Specifically, we (a) described policies that admitted applicants described in terms of specific values (e.g., 1350 SAT score) or a range of values (e.g., 1250+ SAT scores). We also (b) varied whether the background information about the applicants concerned their family's household income or high school's graduation rate. We preregistered this study on AsPredicted.org (https://aspredicted.org/NNO_ZMY).

Methods

Participants. We recruited participants on Amazon Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition or who completed the survey in under 90 seconds. We ended with a final sample of 603 participants (M_{age} = 36.50, SD = 12.34; 51.91% female).

Procedure and materials. Participants were randomly assigned to a cell from a 3 (Type of Decision: Individual, Policy-Specific, Policy-Range) x 2 (Information: Household Income, Graduation Rate) design. As in Study 1, participants making an Individual decision selected one of two individual applicants to admit, whereas participants making a Policy-Specific or Policy-Range decision selected one of two admissions policies. The difference between the Policy-Specific and the Policy-Range decisions is in the description of the applicants affected by the policies.

Before making their decision (about individuals or policies), participants examined a table of information with the two options they could choose from. The tables used in the Individual and Policy-Specific conditions were nearly identical to those used in Study 1 of the

main manuscript. The table used in the Policy-Range condition was identical to the Policy-Specific's table, except that the information in the table represented applicant's attributes in terms of ranges rather than specific values.

As in all studies, the two options in the table described a higher-scoring, advantaged option and a lower-scoring, disadvantaged option. We counterbalanced whether the higher-scoring, advantaged option or the lower-scoring, disadvantaged option was labeled "A" and appeared on the left side of the table. The two options were identical in all respects except for their SAT scores and background information. The background information was either about their "Estimated Household Income" (as in Study 1) or "High School's Graduation Rate". One option had higher scores and a more advantaged background (higher household income or graduation rate), while the other option had lower scores and a more disadvantaged background (lower household income or graduation rate)[5].

In this study, we provided more detail about the individual applicants and policies than in other studies. In the Individual decision, participants also learned where each applicant was from (Minneapolis, MN or Cincinnati, OH) and their intended major (computer science or engineering). We counterbalanced which city and intended major (4 combinations were possible) was associated with higher-scoring, advantaged option or the lower-scoring, disadvantaged option. In the Policy-Specific and Policy-Range decision, participants also read an introductory statement about each policy. The policy option that would favor higher-scoring, advantaged type of applicants was described as "a Results-focused policy. It would admit individuals based exclusively on the attributes that signal ability and potential and are directly comparable across

[5] Tables S2 and S3 report the choice collapsed across the type of information participants saw (household income or high school graduation rate). We found significant differences between the Individual vs. Policy-specific and Individual vs. Policy-range conditions for both types of information for both Choice and Fair Choice (all *p*s < .001).

applicants (e.g., SAT scores, GPA)." The policy option that would favor lower-scoring, disadvantaged type of applicants was described as "a Context-focused policy. It would admit individuals based on the attributes that signal ability and potential and are directly comparable across applicants (e.g., SAT scores, GPA), but it would give more weight to achievements that occurred in the face of adversity."

After reading the information about the options (applicants or policies) and examining the table of information, all participants indicated which of the two options they would select and which option would be the fairer one to select (forced choice: A, B). Participants then explained, in a free-response format, why they chose the particular applicant or policy that they did, as well as what they thought was the difference between the two applicants or policies.

Next, participants indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with two moralization and two objective information items adapted from Kreps and Monin (2014). The two moralization statements were: "Morality is irrelevant to my choice of [applicant/policy] (reverse-scored)" and "My choice of [applicant/policy] is tied to my core moral values and beliefs". The two objective information statements are: "My choice of [applicant/policy] is based on objective criteria" and "My reasons for choosing this [applicant/policy] are concrete and factual". Per the criteria described in our preregistration, we created composite measures of the moralization (label: *Moralization composite*; r = .40) and objective information items (label: *Objective information composite*; *r* = .58).

Next, participants answered five exploratory questions about their perception of the applicants and the decision they faced. Depending on the type of decision participants made, the questions were either about applicants (in the Individual condition) or types of applicant (in the Policy conditions). We describe the measures ahead using the Individual "applicant" version of

the questions. Participants indicated which applicant "…is objectively more qualified for admission?" (label: *Qualified*) and "…faced more adversity in their life?" (label: *Adversity*). For the exploratory questions, participants indicated which applicant's achievements are more "…driven by the role of luck in their lives?" (label: *Luck*) and "…driven by their own effort?" (label: *Effort*). Participants answered these questions using 7-point Likert scales with the options (Applicant A and B) at the ends of the scale (e.g., 1 = *Applicant A is extremely more qualified*, 4 = *The applicants are equally qualified*, 7 = *Applicant B is extremely more qualified*). The last exploratory question asked, "To what extent does selecting a certain applicant imply that the other applicant will or will not be admitted?" (label: *Zero-sum*; 1 = *They will definitely NOT be admitted*, 4 = *They may or may not be admitted*, 7 = *They will definitely be admitted*). While answering all of these questions, participants were able to see the table of information about the applicants or types of applicants. Finally, participants provided demographic information (age, sex, political attitudes, and party affiliation). SM 3 contains all results for non-choice and non-fairness measures for this study.

SM 2.4. STUDY S4 (MANIPULATING FEATURES OF POLICIES)

In Study S4, we directly manipulate three features of policy decisions—the framing of the decision as one about policies, the number of individuals affected and the identifiability of those individuals. We preregistered this study on AsPredicted.org (https://aspredicted.org/VWY_SJI).

Methods

Participants. We recruited online participants on Amazon's Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition or who completed the survey in under 90 seconds. Our final sample comprised 807 participants (M_{age} = 36.32, SD = 12.27; 44.86% female).

Procedure and materials. We randomly assigned participants to one of four conditions (Decision Type: Individual, Policy Narrow, Policy Broad, and Policy Abstract) in a between-subjects design. As in Study 1, we asked all participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. In all conditions, we presented participants with the same table of information as in Study 1, and additionally we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left. Participants in the Individual condition learned that they would be making a choice between "two applicants" and then saw a table of information describing the two specific applicants.

In the Policy-Narrow condition, participants learned that they would be making a choice between two policies that would admit "specific applicants." Participants were told that their choice would "only apply to these two applicants, NOT to a broader pool of applicants." Because participants might find it odd that the policy would only affect two applicants, we gave some

additional detail to make the policy more believable. We explained that the policy would "only affect one final pair of applicants" because "it is almost the end of admissions season." We also did not want participants to speculate whether the policy they chose would affect more applicants down the line, so we further explained that their "fellow admissions officers will create a new policy next year." Then, participants saw a table of information describing the two specific applicants. In sum, participants in the Policy-Narrow condition made a decision that was functionally the same as the decision between individuals, except that the framing of the decision—choosing policies as opposed to choosing individuals—was different.

In the Policy-Broad condition, participants learned that they would be making a choice between two policies that would admit "specific applicants." Participants were told that their choice would "apply to these two applicants, AND a broader pool of these types of applicants." Hence, as in the Policy-Narrow condition, two individuals affected by the policy decision are still identified, but unlike the Policy-Narrow condition, many individuals are affected. As in the other conditions, participants saw a table of information describing the two specific applicants that would be admitted by their choice of policy.

The Policy-Abstract condition was most similar to the policy condition featured in the other studies. In the Policy-Abstract condition, participants learned that they would be making a choice between two policies that would admit different "types of applicants" and that their choice would "apply to a broader pool of these types of applicants." Then, they saw the table of information describing the two types of applicants. Therefore, the Policy-Abstract condition differed from the Policy-Broad condition in that all of the individuals affected by the policy were abstract, rather than identified individuals.

After reading about each applicant or policy, we asked participants to select one of the two applicants or policies (forced choice: A, B) and to indicate which applicant or policy was the fairer one to select (forced choice: A, B). Then, participants indicated their agreement with several mechanism measures.

Participants indicated their agreement with three statements pertaining to moralization (based on Skitka & Morgan, 2014): "My choice of [applicant/policy/rule] is tied to my core moral values and beliefs," "There are very important ethical aspects to this choice," and "My choice of [applicant/policy/rule] is connected to my beliefs about fundamental right and wrong." Participants also responded to two statements pertaining to their reliance on objective information (based on Kreps & Monin, 2014): "My choice of [applicant/policy/rule] is based on objective criteria." and "My reasons for choosing this [applicant/policy/rule] are concrete and factual." Participants indicated their agreement with these statements on a 7-point Likert scale (1 = *Strongly disagree* to 7 = *Strongly agree*).

Participants also reported their perceptions of the importance and consequence of their decision: "How many people will be affected by your decision?" (label: *Scope*; 1 = *No people at all* to 7 = *A large amount of people*), "How important was your decision?" (label: *Importance*; 1 = *Not important at all* to 7 = *Extremely important*), and "How consequential was your decision?" (label: *Consequential*; 1 = *Not consequential at all* to 7 = *Extremely consequential*). A principal axis factor analysis with a Varimax rotation (Russell, 2002) revealed one factor, which omitted the Scope item, so we created a composite (label: *Consequence Composite*) of only the Importance and Consequential items (r = .54). For completeness, we report separate analyses for the Importance and Consequential items in the supplementary materials (SM 3).

At the end of the survey, participants provided an open-ended answer for why they chose the particular applicant or policy that they did and reported demographic information (age, gender, political attitudes, ethnicity, education level, and income level).

Results

We were primarily interested in three contrasts across measures: Individual vs. Policy-Narrow, Policy-Narrow vs. Policy-Broad, and Policy-Broad vs. Policy-Abstract. In order to adjust for multiple comparisons within each measure, we consider the critical significance level to be .017 (i.e., .05/3). For completeness, we report additional contrasts (e.g., Individual vs. Policy-Broad), but we encourage readers to use a more conservative threshold for interpreting those results.

Choice. We found a significant overall effect of Decision Type on participants' choices (χ^2 [3, 807] = 24.97, $p < .001$). Participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (64.73%) than participants choosing between policies in the Policy-Narrow (51.26%; χ^2 [1, 406] = 7.57, $p = .006$), Policy-Broad (48.24%; χ^2 [1, 406] = 11.24, $p = .001$), and Policy-Abstract conditions (40.59%; χ^2 [1, 409] = 23.91, $p < .001$). Notably, even when the number of people affected and the identifiability of those people was held constant (in the Policy-Narrow condition), participants chose the policy that would admit the higher-scoring, higher-income applicant at a lower rate than participants choosing between the applicants. Participants in the Policy-Narrow and Policy-Broad conditions did not choose the higher-scoring, higher-income option at significantly different rates (χ^2 [1, 398] = 0.36, $p = .547$), nor did participants in the Policy-Broad and Policy-Abstract conditions (χ^2 [1, 401] = 2.37, $p = .123$). These last two contrasts reveal that the greater number of individuals affected by the policy decision and the decreased identifiability of those individuals

did not shift participants' choices as much as reframing the individual decision as one about policies.

Fair choice. These judgments followed a similar pattern as Choice, except that the Policy-Broad and Policy-Abstract conditions were significantly different. We found a significant overall effect of Decision Type of participants' fair choice judgment (χ^2 [3, 807] = 56.00, p < .001). Participants choosing between individuals judged the higher-scoring, higher-income option to be the fair choice at a higher rate (71.01%) than participants choosing between policies in the Policy-Narrow condition (55.28%; χ^2 [1, 406] = 10.82, p = .001), Policy-Broad (48.74%; χ^2 [1, 406] = 20.99, p < .001), and Policy-Abstract condition (34.65%; χ^2 [1, 409] = 54.27, p < .001). Participants in the Policy-Narrow and Policy-Broad conditions did not judge the higher-scoring, higher-income option to be the fair one at significantly different rates (χ^2 [1, 398] = 1.70, p = .192), but, in contrast to the choice results, participants in the Policy-Broad and Policy-Abstract conditions did (χ^2 [1, 401] = 8.19, p = .004).

Consequence composite. We found a significant overall effect of Decision Type on the Consequence composite measure, F(3, 803) = 3.66, p = .012. However, none of the key pairwise comparisons were significantly different at p < .017 (Individual vs. Policy-Narrow. p = .308; Policy-Narrow vs. Policy-Broad, p = .060; Policy-Broad vs. Policy-Abstract, p = .197). The remaining contrasts also do not suggest a meaningful difference between conditions (Individual vs. Policy-Broad, p = .377; Individual vs. Policy-Abstract, p = .029; Policy-Narrow vs. Policy-Abstract, p = .002). Hence, it seems that participants thought their decisions were just as consequential for all decision types.

Scope. We found a significant overall effect of Decision Type on the scope measure, F(3, 803) = 117.78, p < .001. Consistent with the intent of our manipulation, participants choosing

between individuals perceived their choice as affecting fewer people than participants choosing abstract policies ($p < .001$) and broad policies ($p < .001$), but not than participants choosing narrow policies ($p = .665$). Participants choosing narrow policies also saw their choice as affecting fewer people than participants choosing abstract policies ($p < .001$) and broad policies ($p < .001$). Participants choosing abstract policies did not perceive their choice to affect a different number of people than participants choosing broad policies ($p = .075$).

Discussion

Participants who made decisions about individuals favored the higher-scoring, higher-income applicant at a higher rate than participants who made decisions about policies, even when those policies affected the same number of individuals and those individuals were identifiable. We considered three features of policy decisions—the high number of individuals affected, the low identifiability of those individuals, and the framing of a decision as one about policies—that could lead people to admit more lower-achieving, disadvantaged applicants. We found that the framing of the decision as one about policies had the largest effect on participants' choices. Because policies are generally seen as statements of a decision maker's or organization's aims, the simple framing of the decision as one about policies may have been sufficient to cue the types of thoughts associated with policy making. The high number of individuals affected and their low identifiability also influenced choice, but the extent to which each of those did so appears to be smaller than the effect of the policy frame. In other words, the step-wise pattern of choice suggests that all three of the policy features we tested may influence choice, but that the policy frame influences choice the most.

SM 2.5. STUDY S5 (MANIPULATING FEATURES OF POLICIES IN HIRING)

In Study S5, as in Study S4, we examined whether people's policy choices change when we manipulate the number of individuals affected by the policies and the identifiability of those individuals. To do so, we tested decisions about policies that affect two specific individuals versus policies that affect many abstract individuals in the context of hiring decisions. Unlike in Study S4, we did not include a policy decision that affects two identified individuals and many other individuals (i.e., the "Policy-Broad" version) because it did not significantly vary from the two others policy versions we tested in Study S4. Additionally, in this study we sought to learn more about what psychologically distinguishes choosing policies from choosing individuals. We preregistered this study on AsPredicted.org (https://aspredicted.org/MMI_YZJ and https://aspredicted.org/ZXN_SKA).

Methods

Participants. We recruited US- and UK-based online participants on Prolific Academic for a small monetary sum. As indicated in our preregistration, we excluded participants who either twice failed an attention check before being assigned to a condition, did not pass a display compatibility check, indicated that the images in the survey did not display correctly, or completed the survey in under 90 seconds. Our final sample comprised 859 participants (M_{age} = 34.19, *SD* = 12.18; 63.80% female).

Procedure and materials. We randomly assigned participants to one of three conditions (Decision Type: Individual, Policy-Narrow, and Policy-Abstract) in a between-subjects design. As in Study 2, we asked all participants to imagine that they were the head of Human Resources at a technology start-up and that they had to decide which job candidates to interview for their start-up's software engineer position. In all conditions, we presented participants with the same

table of information as in Study 2 and we counterbalanced whether the higher-scoring, male option or the lower-scoring, female option was labeled "A" and appeared on the left.

The Individual condition was nearly identical to the Individual condition we tested in Study 2. We told participants that "it is almost the end of the hiring season, and you are making a decision about one pair of candidates. You and your fellow hiring officers will interview new candidates next year." Then, participants saw a table of information "with details about the two specific candidates" that they had to choose between.

In the Policy-Narrow condition, participants learned that they would be making a choice between two policies that would lead to interviewing "two specific candidates." In order to enhance the believability of these policies, which might seem odd because they only lead to interviewing a specific candidate, we explained that "it is almost the end of the hiring season, so your policy will only affect one final pair of candidates. You and your fellow hiring officers will create a new policy next year." Then, participants saw a table of information "with details about the two specific candidates that would be interviewed under each policy." As in Study S4, the Policy-Narrow condition here narrows the number of individuals affected to just two—the same as in the Individual condition—and it makes those individuals identifiable. Hence, the choice of narrow policies is functionally the same as a choice between two individuals, except that it is framed as a choice between policies.

In the Policy-Abstract condition, participants learned that they would be making a choice between two policies that would lead to interviewing different "types of candidates." We told participants that, "it is the beginning of the hiring season, so your policy will affect this year's entire pool of candidates. You and your fellow hiring officers may or may not create a new policy next year." This detail was meant to balance the amount of information conveyed across

conditions. Then, participants saw a table of information "with details about the two types of candidates that would be interviewed under each policy."

After reading about each candidate or policy, we asked participants to select one of the two candidates or policies (forced choice: A, B) and to indicate which candidate or policy was the fairer one to select (forced choice: A, B).

We also measured the extent to which participants made the choice based on values or objective information. Participants first reported how much they relied on their values versus objective information ("Is your choice of [candidate/policy] based more on your values or the objective achievements of the candidates?") on a 7-point Likert scale (1 = *Much more my values*, 4 = *Equally*, 7 = *Much more the details of each candidate*). Then, participants indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with four statements about their macrojustice values and objective information. The two (macrojustice) values statements are: "My choice of [candidate/policy] was motivated by a desire to increase diversity in my organization" and "My choice of [candidate/policy] was motivated by a desire to create equality in my organization."[6] The two objective information statements are: "My choice of [candidate/policy] is based on objective criteria" and "My reasons for choosing this [candidate/policy] are concrete and factual." We combined the two values items (label: *Macrojustice values*; r = .81) and the two objective information items (label: *Objective information composite*; r = .62).

[6]We originally conceptualized of these items as reflecting specific values prioritized when making these decisions, as opposed to relying on objective information to make the decisions. This reflected our earlier theorizing about what might explain the policy-people gap. However, we find limited support that the gap was explained by a differential reliance on values-versus-objective information. Because the values items are very similar to the macrojustice items we collected in Study 4 (as well as Studies S7 and S11), we present them as macrojustice.

We presented two manipulation checks: "How many candidates did you think would be affected by your choice?" (0, 1-2, More than 2) and "In your opinion, which of the following best describes who your choice would affect?" (Two specific candidates, Two types of candidates). Additionally, to verify our theorizing that the framing of a decision as one about policies makes people see it as guiding future decisions, we asked participants to indicate their agreement (1 = *Strongly Disagree*, 4 = *Neither Agree Nor Disagree*, 7 = *Strongly Agree*) with the following statement: "The purpose of selecting a [candidate/policy] is to guide future interviewing decisions my company makes" (label: *Guide*).

Situational attributions. We also asked participants about their attributions for the scores of female candidates and for their confidence in those attributions. It is possible that people making policy decisions might see people's achievements more as the product of factors outside their control (e.g., having an advantage or disadvantage in hiring as a consequence of being male or female) than of factors within their control (e.g., effort, ability) because they think about groups rather than specific individuals (see, e.g., Critcher & Dunning, 2013, and O'Laughlin & Malle, 2002). We asked about attributions both for the scores of female candidates *in general* and for the scores of the *specific* female candidate depicted in the table. The general attribution questions were: "To what extent do you think the Coding scores of women are lower than the Coding scores of men because women are given less opportunities to learn coding?" and "How confident are you that the Coding scores of women are lower than the Coding scores of men because women are given less opportunities to learn coding?". The *specific* version of these questions varied by condition; in the Individual and Policy-Narrow conditions they referred to a "candidate" and in the Policy-Abstract condition to a "candidate type." We labeled the female candidate (or candidate type) based on her position in the table (e.g., "Candidate A," "Candidate

Type B"), which was presented above the questions. For the sake of brevity, we illustrate the phrasing of the questions using the language of the Individual condition. The specific attribution questions were: "To what extent do you think the Coding score of Candidate A is lower than the Coding score of Candidate B because Candidate A was given less opportunities to learn coding, due to their gender?" and "How confident are you that the Coding score of Candidate A is lower than the Coding score of Candidate B because Candidate A was given less opportunities to learn coding, due to their gender?". All questions used a 7-point Likert scale (1 = *Not at all*, 4 = *Moderately*, 7 = *Extremely*) and we counterbalanced whether participants answered the specific or general questions first. A principal axis factor analysis with a Promax rotation (Russell, 2002) suggested one factor explained most of the variance from these four items, so we combined them into one composite measure (label: *Situational Attribution*; $\alpha = .89$).

Exploratory items. We collected two exploratory items about the perceived purpose of the decision. Participants indicated their agreement (1 = *Strongly Disagree*, 4 = *Neither Agree Nor Disagree*, 7 = *Strongly Agree*) with two statements about the purpose of the decision they made: "The purpose of selecting a [candidate/policy] is to make a statement about what I care about." (label: *Statement of Self*) and "The purpose of selecting a [candidate/policy] is to make a statement about what my company values." (label: *Statement of Company*)."

Individual differences. We also measured participants' social-dominance orientation (SDO). Social dominance orientation captures individual differences for group-based hierarchy and inequality (Ho et al., 2015). We measured participants' social-dominance orientation using the SDO_7, a 16-item scale (Ho et al., 2015). The SDO_7 scale comprises two dimensions, dominance (SDO-D) and egalitarianism (SDO-E). We were particularly interested in the egalitarianism dimension because it predicts support for policies that would bring about equality

(Ho et al., 2015), and also predicts perceiving harm and feeling empathy for (dis)advantaged groups (Lucas & Kteily, 2018). It also influences the attributions that people draw for the success and failure attributions of (dis)advantaged individuals (McClanahan, Ho, & Kteily, 2019). We expected that individual differences for the egalitarian dimension might moderate the effect of Decision Type on Choice, but we did not find evidence of this. We report the results in SM 4.

At the end of the survey, participants provided an open-ended answer about why they chose the particular candidate/policy that they did, their employment industry and hiring experience (if any), and demographic information (age, gender, and political attitudes).

Results

We were primarily interested in two contrasts across measures (excluding the manipulation check): Individual vs. Policy-Narrow and Policy-Narrow vs. Policy-Abstract. In order to adjust for multiple comparisons within each measure, we consider the critical significance level to be .025 (i.e., .05/2). For completeness, we report the Individual vs. Policy-Abstract contrast as well, but we encourage readers to use a more conservative threshold for interpreting those results.

Manipulation checks. The majority of participants in each condition answered the manipulation check questions as we expected. In the Individual condition, 76.47% of participants indicated that their choice would affect 1-2 candidates and 69.90% indicated that the candidates could be best described as specific candidates. In the Policy-Narrow condition, 70.32% of participants indicated that their choice would affect 1-2 candidates and 58.30% indicated that the candidates could be best described as specific candidates. Finally, in the Policy-Abstract condition, 59.58% indicated more than *2* candidates and 62.37% indicated that the candidates could be best described as types of candidates. Less than 5% of participants in each condition

indicated that their choice would affect no candidates. In sum, the majority of participants in each condition perceived the number of individuals affected by their choice, and the identifiability of those individuals, as intended.

We found a significant overall effect of Decision Type on the Guide measure, $F(2, 856) = 55.42$, $p < .001$. Participants choosing between individuals were less likely to believe that the purpose of their choice was to guide their company's future choices ($M = 4.81$, $SD = 1.52$), relative to participants choosing between narrow policies ($M = 5.66$, $SD = 1.15$; $p < .001$) or abstract policies ($M = 5.83$, $SD = 1.01$; $p < .001$). Participants choosing between narrow policies and abstract policies did not significantly differ in their agreement ($p = .115$). In other words, the Policy-Narrow and Individual decisions were framed as having equivalent consequences, yet participants still believed this policy decision was intended to guide future decisions more than the individual decision. This result is consistent with our theorizing that people perceive policies as general guides for decisions.

Choice. We found a significant overall effect of Decision Type on participants' choices (χ^2 [2, 859] = 10.10, $p = .006$). Participants choosing between individuals did not choose the higher-scoring, male option at a significantly higher rate (73.70%) than participants choosing between policies in the Policy-Narrow condition (72.79%; χ^2 [1, 572] = 0.06, $p = .806$). Participants choosing individuals and narrow policies did choose the higher-scoring, male option at significantly higher rates than participants choosing abstract policies (62.72%; Individual vs. Policy-Abstract: χ^2 [1, 576] = 8.02, $p = .005$; Policy-Narrow vs. Policy-Abstract: χ^2 [1, 570] = 6.62, $p = .010$).

Fair choice. As with Choice, we found a significant overall effect of Decision Type of participants' fair choice judgment (χ^2 [2, 859] = 27.04, $p < .001$). In contrast to the choice results,

participants choosing between individuals judged the higher-scoring, male option to be the fair choice at a higher rate (84.43%) than participants choosing between policies in the Policy-Narrow condition (76.33%; χ^2 [1, 572] = 5.96, p = .015). Participants in the Individual and Policy-Narrow conditions judged the higher-scoring, male option to be the fair one at significantly different rates than participants in the Policy-Abstract condition (65.85%; Individual vs. Policy-Abstract; χ^2 [1, 576] = 26.62, p < .001; Policy-Narrow vs. Policy-Abstract: χ^2 [1, 570] = 7.60, p = .006).

Values. We found a marginally significant overall effect of Decision Type, F(2, 858) = 3.58, p = .028. Participants choosing between individuals relied less on their values than participants abstract policies (p = .008), but not narrow policies (p = .142). Participants choosing between narrow policies and abstract policies did not differ (p = .230).

Objective information. We found a significant overall effect of Decision Type, F(2, 858) = 7.36, p < .001. Participants choosing between individuals relied more on objective information than participants choosing between narrow policies (p = .006) and abstract policies (p < .001). Participants choosing between narrow policies and abstract policies did not differ (p = .338).

Situational attribution. We found a marginally significant overall effect of Decision Type on the Situational Attribution measure, F(2, 856) = 2.85, p = .058, but the pairwise comparisons of interest were not significantly different: Individual vs. Policy-Narrow (p = .474) and Policy-Narrow vs. Policy-Abstract (p = .109). However, the contrast between the Individual and Policy-Abstract condition suggests a potential difference (p = .020). Overall, though, we do not conclude that people making policy decisions make different attributions about candidates' achievements than people making individual decisions.

Statement of self. We found a significant overall effect of Decision Type on the Statement of Self measure, $F(2, 856) = 8.39$, $p < .001$. Participants choosing between individuals were less likely to believe that the purpose of their choice was to make a statement about what they cared about, relative to participants choosing between narrow policies ($p = .001$) or abstract policies ($p < .001$). Participants choosing between narrow policies and abstract policies did not differ in their agreement ($p = .838$).

Statement of company. We found a significant overall effect of Decision Type on the Statement of Company measure, $F(2, 856) = 14.07$, $p < .001$. Participants choosing between individuals were less likely to believe that the purpose of their choice was to make a statement about what their company cares about, relative to participants choosing between narrow policies ($p < .001$) or abstract policies ($p < .001$). Again, participants choosing between narrow policies and abstract policies did not differ in their agreement ($p = .902$).

Discussion

Participants who made decisions about individuals selected a higher-scoring, male candidate, over a lower-scoring, female candidate, at a higher rate than participants who selected "abstract" policies that would favor that type of candidate. However, participants who faced a choice between "narrow" policies that would specifically select those candidates, and no one else, did not choose differently than participants selecting the specific candidates. However, more participants choosing narrow policies than participants choosing between specific individuals thought that the fair option to choose would be the one that favored the female candidate. It is possible that this discrepancy between choice and fairness reflects differences between the contexts of college admissions and hiring. People may find it less acceptable to *not* hire the most qualified job candidates than to *not* admit the most qualified college applicants. In

other words, the norms of the hiring context may constrain people's choices more than the norms of the admissions context, which may be why in this study we do not see a divergence between people's choices, but we do for what they think is fair to choose. Consistent with this logic, we also find stronger effects on fairness than on choice in Study 3, the other study in the context of hiring. Lastly, we also found that people making policy decisions rely more on their values and less on objective criteria, and that this is associated with favoring lower-achieving, disadvantaged people.

SM 2.6. STUDY S6 (POLICY VS. RULE)

In Study S6, we tested whether the individual-versus-policy divergence is driven by the word "policy" itself, or rather by the construct of a *policy* (i.e., a guide for future decisions). It is possible that the word "policy" might cue certain associations, such as the policy positions of one's political party, which might account for the gap. To test this possibility, we used the word "rule" instead of "policy" to describe the policy decision frame. We chose the word "rule" because it has a similar meaning as "policy" but is likely less politicized. Additionally, we measured our hypothesized mechanism, namely, how much people rely on their values versus objective information when making each type of decision. We preregistered this study on AsPredicted.org (https://aspredicted.org/QMB_LNE).

Methods

Participants. We recruited online participants on Prolific Academic for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition. Our final sample comprised 593 participants (M_{age} = 34.96, *SD* = 11.19; 48.74% female).

Procedure and materials. We randomly assigned participants to one of three conditions (Type of Decision: Individual, Policy, Rule). This design replicates the procedure of Study 1 with two key modifications. First, the design is fully between subjects. Using the same table of information as in Study 1, participants made only one type of decision, rather than both. As in Study 1, the table described the applicants (in the Individual condition) or types of applicants (in the Policy and Rule conditions). Second, this design includes a new type of decision. In the Rule type of decision, we replaced all instances of the word "policy" with "rule." In sum, participants chose applicants for admission, admissions policies, or admissions rules.

After participants read about each applicant, policy, or rule, we asked them to select one of the two options (forced choice: A, B) and to indicate which option was the fairer one to select (forced choice: A, B). Next, participants indicated their agreement with several potential mechanism measures.

Participants indicated their agreement with three statements pertaining to moralization (based on Skitka & Morgan, 2014): "My choice of [applicant/policy/rule] is tied to my core moral values and beliefs.", "There are very important ethical aspects to this choice.", and "My choice of [applicant/policy/rule] is connected to my beliefs about fundamental right and wrong." Participants also responded to two statements pertaining to their reliance on objective information (based on Kreps & Monin, 2014): "My choice of [applicant/policy/rule] is based on objective criteria." and "My reasons for choosing this [applicant/policy/rule] are concrete and factual." Participants indicated their agreement with these statements on a 7-point Likert scale (1 = *Strongly disagree* to 7 = *Strongly agree*). Participants also answered a related question: "Is your choice of [applicant/admissions policy/admissions rule] based more on your values or the objective achievements of [each applicant/the applicant associated with the policy/the applicant associated with the rule]?" (label: *Values-versus-objective-information*; 1 = *Much more my values*, 4 = *Equally my values and the details of each applicant*, 7 = *Much more the details of each applicant*;). Last, participants indicated their agreement with the following statement: "My choice of [applicant/policy/rule] makes a statement about what is important to me (label: *Statement;* 1 = *Strongly disagree* to 7 = *Strongly agree*)." We report analyses for a composite of the moralization items (label: *Moralization composite*; $\alpha = .83$), a composite of the objective information items (label: *Objective composite*; $r = .59$), the single Values vs. Achievement item, and the last Statement item in the supplementary materials (SM 3).

At the end of the survey, participants provided an open-ended answer for why they made their choice of applicant/policy/rule and demographic information (age, gender, political attitudes).

Results

Choice. We found a significant overall effect of Decision Type on participants' choices (χ^2 [2, 593] = 40.21, $p < .001$). Participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (63.13%) than participants choosing between policies (36.41%; χ^2 [1, 382] = 27.23, $p < .001$) or rules (35.07%; χ^2 [1, 409] = 32.19, $p < .001$). Participants choosing between policies and rules did not choose the higher-scoring, higher-income option at significantly different rates (χ^2 [1, 395] = 0.08, $p = .781$).

Fair choice. As with Choice, we found a significant overall effect of Decision Type on participants' fair choice judgment (χ^2 [2, 593] = 93.59, $p < .001$). Participants choosing between individuals judged the higher-scoring, higher-income option to be the fair choice at a significantly higher rate (70.20%) than participants choosing between policies (32.07%; χ^2 [1, 382] = 55.56, $p < .001$) or rules (26.07%; χ^2 [1, 409] = 79.80, $p < .001$). Participants choosing between policies and rules did not judge the higher-scoring, higher-income option to be the fair option at significantly different rates (χ^2 [1, 395] = 1.72, $p = .189$).

Moralization. We conducted a one-way ANOVA on the Moralization composite and found a significant effect of Decision Type, $F(2, 590) = 19.35$, $p < .001$. Participants choosing between individuals saw their choice as less moralized ($M = 4.39$, $SD = 1.30$) than participants choosing between policies ($M = 5.11$, $SD = 1.27$; $p < .001$) and rules did ($M = 5.05$, $SD = 1.23$; $p < .001$)

Objective Information. We conducted a one-way ANOVA on the Objective Information composite and found a significant effect of Decision Type, $F(2, 590) = 4.28$, $p = .014$. Participants choosing between individuals saw their choice as relying more on objective information ($M = 5.50$, $SD = 1.21$) than participants choosing between policies ($M = 5.26$, $SD = 1.11$; $p = .045$) and rules did ($M = 5.17$, $SD = 1.14$; $p = .006$)

Discussion

Participants who made decisions about rules favored the lower-scoring, lower-income applicants at the same rate as participants who made decisions about policies. Hence, the gap between individual and policy decisions seems to be caused by more than just associations cued by the word policy.

SM 2.7. STUDY S7 (SHIFTING STANDARDS OF FAIRNESS)

In Study S7, as in Study 3, we directly test the mechanism proposed by our theoretical account. We test whether the gap between policy and individual decisions is explained by decision makers prioritizing macrojustice principles and deprioritizing microjustice principles when they make policy decisions, relative to individual decisions (Hypothesis 2). We also compare the individual and policy decisions in a more controlled way in this study. Specifically, we compare an individual admissions decision (choosing one of two applicants) to a policy decision that would only affect two applicants. We preregistered this study on AsPredicted.org (https://aspredicted.org/TCC_VYN).

Methods

Participants. We recruited US-based online participants on Prolific Academic for a small monetary sum. Using Prolific's filters, we recruited participants with at least a 4-year college degree to ensure that they had some familiarity with the college admissions context; 95.5% of participants reported having at least a bachelor's degree. As indicated in our preregistration, we excluded participants who failed an attention check or who answered the three questions about macrojustice and microjustice (described below) in under 6 seconds.[7] Our final sample comprised 283 participants (M_{age} = 33.04, *SD* = 10.17; 49.12% female).

Procedure and materials. We randomly assigned participants to one of two conditions (Decision Type: Individual, Policy Framing) in a between-subjects design. As in Studies 1 and 2, we asked all participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. In all conditions, we

[7] For robustness we also conducted our analyses including the 15 participants (5% of the sample) who answered the questions in under 6 seconds. All results are practically identical. For example, in the Control condition, participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (57.14% with these participants vs. 56.77% without them) than participants choosing between policies (43.80% vs. 43.75%; p = .022 vs. p = .029).

presented participants with the same table of information as in Study 1 and we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left.

Participants in the Policy Framing condition learned that they would be making a choice between two policies that "only affect one final pair of applicants this year." Because participants might find it odd that the policy would only affect two applicants, we gave some additional detail to make the policy more believable. We explained that the policy would "only affect one final pair of applicants" because "it is almost the end of admissions season." We also did not want participants to speculate whether the policy they chose would affect more applicants down the line, so we further explained that they and their "fellow admissions officers can create a new policy next year." Then, participants saw a table of information describing "the two specific applicants that would be admitted under each policy."

Participants in the Individual condition learned the same information (that "it is almost the end of admissions season" and that they were making a decision about "one final pair of applicants") but then simply saw a table of information describing the two specific applicants and made a choice about whom to select. In sum, participants in the Individual and Policy Framing conditions made functionally equivalent decisions, except that the framing of the decision—choosing policies versus choosing individuals—was different.

After reading about each applicant or policy, participants selected one of the two applicants or policies (forced choice: A, B), indicated which applicant or policy was the fairer one to select (forced choice: A, B), and reported how likely they were to pick either option (1 = *Much more likely to choose A,* 4 = *Equally Likely*, 7 = *Much more likely to choose B*). Participants also reported how much their decision was motivated by concerns related to the

macrojustice and microjustice. Specifically, participants reported how much their decision was motivated (1 = *Not at all*, *4 = Moderately*, 7 = *Extremely*) by "creating equality" (macrojustice), "increasing diversity" (macrojustice), and "rewarding applicants for their accomplishments" (microjustice). These items are similar to those tested in Study 3. As indicated in our preregistration, we averaged the two macro items into a *Macrojustice* composite because they were strongly correlated (r = .76), but we did not combine all three items for our main analysis because their reliability was lower than our threshold (α = .63). Participants also provided demographic information (age, gender, race/ethnicity, level of education, income, and political attitudes).

Results

Choice. Consistent with Hypothesis 1, participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (56.77%) than participants choosing between policies with equivalent consequences (43.75%; χ^2 [1, 283] = 4.76, p = .029).

Likelihood. As with Choice, participants reported a greater likelihood of choosing the higher-scoring, higher-income option when choosing between individuals (M = 4.34, SD = 1.96) than between equivalent policies (M = 3.68, SD = 1.88; t[281] = 2.85, p = .005).

Fair choice. Consistent with our theoretical account, participants also indicated that it was fair to choose the higher-scoring, higher-income option at a significantly higher rate when choosing between individuals (63.23%) than choosing between equivalent policies (32.81%; χ^2 [1, 283] = 25.94, p < .001).

Mediation analyses. We also tested whether the data are consistent with the proposed causal model in which the effect of making an individual versus policy decision on choosing the higher-scoring, higher-income option is mediated by one's macrojustice and microjustice

concerns (Hypothesis 2; see manuscript's Figure 1). Assuming this causal model (the same one we tested in Study 3), we conducted binary mediation analysis setting the choice of applicant/policy as the dependent variable (1 = higher-scoring, higher-income option, 0 = lower-scoring, lower-income option), the type of decision as the independent variable (1 = Individual, 0 = Policy Framing), and the macrojustice composite and the single microjustice item as simultaneous mediators. We ran the model using the PROCESS Macro for SPSS (Model 4, Hayes, 2018), which allowed us to estimate the indirect effect using a percentile bootstrap estimation approach with 10,000 samples.

We found that the participants making the Individual decision, compared to the Policy Framing decision, were motivated significantly less by macrojustice concerns (b = -0.75, p = .001), which was significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -0.96, p < .001; macrojustice indirect effect = 0.72, 95% *CI* = [0.28, 1.26]). However, participants making the Individual decision, compared to the Policy Framing one, were not significantly more (or less) motivated by microjustice concerns (b = 0.10, p = .578); though the microjustice concerns were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option (b = 0.76, p < .001; microjustice indirect effect = 0.08, 95% *CI* = [-0.20, 0.38]). The indirect effect of macrojustice was strong enough to fully account for the relationship between making an Individual versus Policy decision and decision makers' choices. In other words, the direct effect (= 0.21, p = .543) was not significant when the model included the macrojustice and microjustice concerns (or the macrojustice concern on its own). We found these same pattern of results when we conducted the same mediation analysis but with likelihood as the dependent variable. We also conducted two other

mediations (one for choice and another for likelihood) where we examined only one mediator, a composite of the macrojustice and microjustice items, with the microjustice item reverse-coded.

Additional mediation. Likelihood, simultaneous mediators. Participants making the Individual decision, compared to the Policy Framing decision, were motivated significantly less by macrojustice concerns (b = -0.75, p = .001), which was significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -0.52, p < .001; macrojustice indirect effect = 0.39, 95% *CI* = [0.15, 0.64]). However, participants making the Individual decision, compared to the Policy Framing one, were not significantly more (or less) motivated by microjustice concerns (b = 0.10, p = .578); though the microjustice concerns were significantly associated with a higher likelihood of choosing the higher-scoring, higher-income option (b = 0.29, p < .001; microjustice indirect effect = 0.03, 95% *CI* = [-0.07, 0.14]).

Additional mediation. Choice, single mediator. Participants making the Individual decision, compared to the Policy Framing decision, were significantly less motivated by microjustice-to-macrojustice concerns (b = -0.53, p = .003), which was significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -1.48, p < .001; indirect effect = 0.79, 95% *CI* = [0.27, 1.37]).

Additional mediation. Likelihood, single mediator. Participants making the Individual decision, compared to the Policy Framing decision, were significantly less motivated by microjustice-to-macrojustice concerns (b = -0.53, p = .003), which was significantly associated with a lower likelihood of choosing the higher-scoring, higher-income option (b = -0.79, p < .001; indirect effect = 0.42, 95% *CI* = [0.15, 0.70]).

Discussion

Participants who made decisions about individuals favored the higher-scoring, higher-income applicant at a higher rate than participants who made decisions about policies, even when those policies had identical consequences. Similarly, participants indicated that the higher-scoring, advantaged applicant was the fair applicant to select more often than they indicated that the policy that would favor that applicant was the fair option to select.

We also tested Hypothesis 2 using mediation analysis. We found that decision makers' macrojustice concerns, but not their microjustice ones concerns significantly mediated the relationship between making a policy versus individual decision and their choice of applicant/policy. This result is the converse of Study 3, where we found that the microjustice concerns, but not the macrojustice ones mediated the relationship. Nevertheless, both these results and those of Study 3 are consistent with H2 – that a greater relative reliance on microjustice (versus macrojustice) accounts for the relationship between making an individual versus policy versus and the likelihood of favoring higher-achieving, advantaged individuals over lower-achieving, disadvantaged individuals.

SM 2.8. STUDY S8 (MICROJUSTICE AND MACROJUSTICE DEVELOPMENT)

In Study S8 we sought to test measures of macrojustice and microjustice within our context. We developed measures by adapting existing measures of macrojustice and microjustice used in other contexts. These measures ultimately proved not to be predicted by making an individual versus policy decision, although they did predict choice. In contrast to the measures tested in Studies 3 and S9, these measures did not ask directly whether participants' choice was motivated by certain concerns, but rather how much certain concerns were relevant. In other words, these measures were more akin to attitudinal measures, which participants shared to the same extent regardless of which type of decision they made. We preregistered Study S8 on AsPredicted.org (https://aspredicted.org/79C_T1Q).

Methods

Participants. We recruited US-based online participants on Amazon Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who failed either of two attention checks or failed a compatibility check for their browser. Our final sample comprised 240 participants (M_{age} = 39.52, *SD* = 10.96; 48.33% female).

Procedure and materials. We randomly assigned participants to one of two conditions (Decision Type: Individual, Policy) in a between-subjects design. The procedure for the Individual and Policy conditions was identical to Study 3. After seeing the information about the applicants or policies, participants selected one of the two applicants or policies (forced choice: A, B), indicated which applicant or policy was the fairer one to select (forced choice: A, B), and reported how likely they were to pick either option (1 = Much more likely to choose A, 4 = Equally Likely, 7 = Much more likely to choose B).

Justice items. We measured the extent to which participants thought several justice-related concerns were important. We asked participants, "How important were the following

concerns for choosing an applicant / a policy? (1 *Not at all important* - 7 *Extremely Important*)" for twenty-four items pertaining to microjustice (5 items), macrojustice (5), procedural justice (9), and distributive justice (5). The microjustice items were based on measures from Sinclair & Mark (1991; items 1-2) and Zdaniuk & Bobocel (2011; items 3-5). The macrojustice items were also based on measures from Sinclair & Mark (1991; items 1-3) and Zdaniuk & Bobocel (2011; items 4-5). The procedural justice items comprised seven items adapted from Colquitt (2001), which were based on items from Thibaut & Walker (1975) and Leventhal (1980), and two items adapted from Lind & Tyler (1997). The distributive justice items comprised four items adapted from Colquitt (2001), which were based on items from Leventhal (1976), and one item adapted from Brockner et al. (2000). The items were the following:

- "College admission should be based on how much applicants have accomplished." (*Micro 1)*
- "College admission should be based on how hard applicants work relative to others." (*Micro 2)*
- "The procedures used in college admissions to determine which applicants should be admitted should be the same for everyone regardless of each applicant's circumstance." (*Micro 3)*
- "The procedures used in college admissions to determine which applicants should be admitted should take into consideration only applicants' objective qualifications." (*Micro 4)*
- "College admission should be given to the applicant with the greatest merit." (*Micro 5)*
- "College admission should be made possible for anyone who is interested." (*Macro 1)*
- "College admission should be equally likely for applicants of all socioeconomic groups." (*Macro 2)*
- "There is too much of a difference in educational opportunities given to disadvantaged and advantaged applicants." (*Macro 3)*
- "There should be an appropriate balance kept between the number of admissions spots given to relatively disadvantaged and advantaged applicants." (*Macro 4)*

- “The number of admission spots given to disadvantaged applicants should not be too much less than the number of spots given to advantaged applicants.” (*Macro 5)*
- “The degree to which applicants are able to express their views and feelings during the admissions process.” (*Proc. 1)*
- "The degree to which applicants have influence over the admissions decision arrived at through the admissions process.” (*Proc. 2)*
- “The degree to which the admission processes is applied consistently.” (*Proc. 3)*
- “The degree to which the admission processes is free of bias.” (*Proc. 4)*
- “The degree to which the admissions process is based on accurate information.” (*Proc. 5)*
- “The degree to which applicants are able to appeal the admissions decision arrived at through the admissions process.” (*Proc. 6)*
- "The degree to which the admissions process upholds ethical and moral standards.” (*Proc. 7)*
- “The degree to which the methods used were very fair.” (*Proc. 8)*
- “The degree to which the methods used to make an admissions decisions were equally fair to both applicants.” (*Proc. 9)*
- “The degree to which admissions decisions reflect the effort the applicants have put into their academic work.” (*Dist. 1)*
- “The degree to which admissions decisions are appropriate given the academic work the applicants have completed.” (*Dist. 2)*
- “The degree to which admissions decisions reflect what the applicants have achieved as students.” (*Dist. 3)*
- “The degree to which admissions decisions are justified, given the applicants' performances.” (*Dist. 4)*
- “The degree to which the outcome of the admissions decision was favorable.” (*Dist 5)*

For each of the constructs (*Micro concerns, Macro concerns, Procedural*, and *Distributive*) we created scales. We depict the correlations between Decision Type, Choice, and the scales in the table below; we include the reliability of the scales in the table as well. We also asked participants to report how much their choice was motivated by the same items (*Macrojustice* based on 2 items and *Microjustice* based on 1 item) used in Study 4, using the

same 7-point Likert scale. We depict the results of t-tests for each of these scales in SM 3. At the end of the study, participants provided demographic information (age, gender, race/ethnicity, level of education, income, and political attitudes).

Results and Discussion

Choice. Participants choosing between individuals chose the higher-scoring, higher-income option at a significantly higher rate (59.84%) than participants choosing between policies (45.13%; χ^2 [2, 240] = 5.10, p = .023).

Likelihood. As with Choice, participants reported a greater likelihood of choosing the higher-scoring, higher-income option when choosing between individuals (M = 4.24, SD = 2.06) than between equivalent policies (M = 3.65, SD = 2.20; t[238] = 2.15, p = .033).

Fair choice. Participants making a decision about individuals were more likely to indicate that the higher-scoring, higher-income applicant (62.99%) was the fair choice than participants making a choice about policies (39.82%; χ^2 [2, 240] = 12.86, p < .001).

TABLE S6
Correlations between Key Independent and Dependent Variables (Study S8)

Variable	1	2	3	4	5	6
1. Decision Type	-					
2. Choice	0.15*	-				
3. Micro (α = .83)	-0.06	0.36***	-			
4. Macro (α = .83)	-0.08	-0.42***	-0.17**	-		
5. Procedural Justice (α = .85)	-0.05	-0.03	-0.34***	0.45***	-	
6. Distributive Justice (α = .85)	-0.06	0.18**	0.64***	0.14*	0.50***	-

Note. Decision Type was coded "1" for the Individual condition and "0" for the Policy condition. Choice was coded "1" when selecting the higher-scoring, advantaged option and "0" when selecting the lower-scoring, disadvantaged option. Alpha denotes reliability for each scale.

† p < .10
* p < .05
** p < .01
*** p < .001

Exploratory factor analyses. As preregistered, we conducted exploratory factor analyses with all items to determine which items to include in macro and micro scales. Maximum likelihood factor analysis (with a cut-off point of 0.50 and the Kaiser's criterion of eigenvalues greater than 1), followed by an oblique (promax) rotation, suggested three or four factors suggesting the same macrojustice and microjustice composites. Two of the factors comprised all of the macrojustice items and the two other factors (or one factor) comprised microjustice items 1-2 and 5. We averaged all the macro items into one scale, which revealed high reliability (Cronbach's α = .83, Coefficient H = .87). We averaged micro items 1-2 and 5 into one scale, which also revealed high reliability (Cronbach's α = .82, Coefficient H = .86).

SM 2.9. STUDY S9 (MICROJUSTICE AND MACROJUSTICE CONTENT VALIDATION)

In Study S9 we sought to develop valid measures of macrojustice and microjustice within our context. We used the Hinkin and Tracey (1999) validation approach and the guidelines proposed by Colquitt, Sabey, Rodell, & Hill (2019). This exercise allowed us to determine "the degree to which a scale's items correspond to the construct's definition" (definitional correspondence) and "the degree to which a scale's items correspond more to the construct's definition than to the definitions of other orbiting constructs" (definitional distinctiveness; Colquitt et al., 2019). Correspondence and distinctiveness help to assess the quality of a new scale. We preregistered Study S9 on AsPredicted.org (https://aspredicted.org/NH6_P1K).

Methods

Participants. We recruited US-based online participants on Amazon Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who failed either of two attention checks, did not complete the study in full, or spent less than 10 seconds rating items for each of the constructs (far less than required to complete the tasks attentively). Our final sample comprised 149 participants (M_{age} = 47.69, *SD* = 10.31; 34.90% female).

Procedure and materials. We asked participants to assess the degree to which survey items matched the definitions of different concepts. Following Colquitt et al.'s (2019) materials, we first familiarized participants with the instructions and had them practice the task with a sample concept and items. After this introduction, participants rated how well each of 20 items matched the definitions of macrojustice, microjustice, procedural justice, and distributive justice. Participants rated the 20 items on four occasions, once for each construct; the order in which we presented the concepts and items was random. The 20 items intended to measure the four concepts: four items for microjustice, four for macrojustice, four for distributive justice (adapted

from Colquitt, 2001, based on Leventhal, 1976), seven for procedural justice (adapted from Colquitt, 2001, based on Thibaut & Walker, 1975, and Leventhal, 1980), and one attention check item. Participants made their ratings on a seven-point Likert scale (1 = *Item is EXTREMELY BAD at measuring the concept* to 7 = *Item is EXTREMELY GOOD at measuring the concept*). We include all items and concept definitions in Table S7. The microjustice and macrojustice items were informed by the original conceptualization of the constructs (Brickman et al., 1981) and items we tested in previous studies (Study S8).

In order to assess the definitional distinctiveness of the macrojustice and microjustice survey items, we selected distributive justice (instantiated as an equity principle; Deutsch, 1975; Leventhal, 1976) and procedural justice as "orbiting constructs." These particular constructs provide a useful benchmark because they are conceptually relevant to macrojustice and microjustice (e.g., they are also standards of fairness that are relevant to selection contexts), and because their definitions are well understood and the scales are used frequently.

Results and Discussion

We calculated Hinkin and Tracey (1999) definitional correspondence and definitional distinctiveness scores for each scale. We interpreted these scores using Colquitt et al.'s (2019) evaluation criteria. The interpretation of the scores depends on the correlation between the scales (when participants actually employ the scales rather than evaluate them) and we do not have these correlations among all the scales, so we provide interpretations assuming a moderate correlation (between .33 and .50, based on Colquitt et al.). Moreover, to situate the quality of these scales, we report the correspondence and distinctiveness summary statistics from Colquitt et al.'s tests of 112 scales published in industrial/organizational psychology and organizational behavior journals between 2010 and 2016.

TABLE S7
Construct Definitions and Item Text (Study S9).

Construct	Item	Definitions and Item Text
Macrojustice		Concern for having a fair overall distribution of opportunities; specifically, how rewards and opportunities are distributed across people in the aggregate.
	1	Making sure that there is a diverse population within a given context
	2	Creating greater equality in the opportunities disadvantaged and advantaged people receive
	3	Balancing the number of opportunities given to disadvantaged and advantaged people
	4	Making sure that there are not too few disadvantaged people in a given context
Microjustice		Concern for what is fair to individuals; specifically, whether individuals' specific qualities determine what opportunities they get.
	1	Selecting people based on their hard work
	2	Selecting people based on their qualifications
	3	Selecting people based on their potential to succeed
	4	Selecting people based on their accomplishments
Distributive Justice		Concern for having fair decision outcomes; specifically, whether rewards and opportunities are allocated in a way that reflects people's contributions.
	1	Making a decision that reflects the effort people have put into their work
	2	Making a decision that is appropriate for the work people have completed
	3	Making a decision that reflects what people have contributed
	4	Being able to justify the outcome, given people's performances
Procedural Justice		Concern for having a fair decision-making processes; specifically, whether processes are consistent, accurate, unbiased, and open to voice and input.
	1	Making sure people are able to express their views and feelings during the decision process
	2	Making sure people have influence over the decision process used to arrive at a decision
	3	Having a consistent decision process
	4	Having a decision process that is free of bias
	5	Having a decision process that is based on accurate information
	6	Making sure people are able to appeal the decision arrived at through the decision process
	7	Having a decision process that upholds ethical and moral standards

Note. The distributive justice items were adapted from Leventhal (1976). The procedural justice items 1-2 were adapted from Thibaut & Walker (1975) and items 3-7 were adapted from Leventhal (1980).

Definitional correspondence. The mean correspondence score for each scale was as follows: macrojustice (*M* = .80, *SD* = .19), microjustice (*M* = .77, *SD* = .19), distributive justice (*M* = .87, *SD* = .15), and procedural justice (*M* = .83, *SD* = .12). The scales examined in Colquitt et al. ranged from .60 to .96, *M* = .87, *SD* = .06. These results reveal that participants judged the macrojustice and microjustice scales to correspond to their construct definitions as well as the procedural justice scale corresponded to its construct and only slightly worse than the distributive justice scale did. The correspondence of our scales was also within the range of scales typically used in management research.

Definitional distinctiveness. The mean distinctiveness score for each scale was as follows: macrojustice (*M* = -.01, *SD* = .07), microjustice (*M* = -.02, *SD* = .07), distributive justice (*M* = .03, *SD* = .06), and procedural justice (*M* = .01, *SD* = .06). As Colquitt et al. (2019) explain, the range for the distinctiveness score is from -1 to 1, where a positive value reflect items receiving higher ratings on the intended construct than on orbiting ones and a negative value reflects items receiving lower ratings on the intended construct than on the orbiting ones.The scales examined in Colquitt et al. ranged from -.04 to .64, *M* = .27, *SD* = .14). These results reveal that the scales were perceived to be very similar to each other because of similarities among items or the constructs' definitions. Because all these scales relate to just ways to make selection decisions, it is possible that participants interpreted them as being related. It is worth noting that the distinctiveness score of each construct's scale is a function not just of the scale's items but also the other scale items that were compared to. Hence, while the distinctiveness of all these scales is low, we believe this greatly reflects the choice of orbiting constructs, as well as the fact that we presented three orbiting constructs for each scale (Colquitt et al., 2019, present

only two orbiting constructs), thus potentially making it more difficult for participants to distinguish between the scales' items.

In sum, while the face validity of the scales we developed were similar to frequently used justice scales. Therefore, they provide a useful first step in measuring the constructs of macrojustice and microjustice and allowing us to investigate their role in policy and individual decisions. Nevertheless, future research should elucidate the relation among all these justice constructs and improve the precision of their corresponding scales.

SM 2.10. STUDY S10 (AN INTERVENTION TO ALIGN CHOICES)

We sought to test a potential intervention to align policy and individual choices. Recent efforts by policy makers and practitioners in college admissions suggest that highlighting adversity may help to align policy and individual decisions, particularly as they pertain to equity (Belkin, 2019). For example, highlighting high school students' contexts led decision makers to admit more low-income applicants (Bastedo & Bowman, 2017). Similarly, we tested an intervention that made salient the adversity that applicants have faced. The intervention consisted of describing the differences between the choice options and, in doing so, highlighting that one (type of) applicant had faced adversity. We predicted that making adversity salient would introduce a new attribute—an applicant's ability to overcome adversity—for which decision makers could give credit to lower-scoring, disadvantaged applicants. We preregistered Study S10 on AsPredicted.org (https://aspredicted.org/FQR_VSM).

Methods

Participants. We recruited online participants on Amazon's Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition or who completed the survey in under 90 seconds. Our final sample comprised 595 participants (M_{age} = 37.30, *SD* = 10.99; 52.27% female).

Procedure and materials. We randomly assigned participants to a condition from a 2 (Decision Type: Individual, Policy) x 2 (Adversity: Salient, Not Salient) between-subjects design. As in Study 1, we asked all participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. In all conditions, we presented participants with the same table of information as in Study 1 and,

additionally, we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left.

The Adversity Not Salient versions of the Individual and Policy conditions were identical to the Individual and Policy conditions we tested in Study 1 (except for the within-subjects portion). In the Adversity Salient conditions, we highlighted information about the applicants, namely whether they faced adversity or not. In doing so, we sought to introduce an attribute—an applicant's ability to overcome adversity—for which decision makers could reward lower-scoring, lower-income applicants, thereby legitimizing the selection of lower-scoring applicants even when broader values were not salient. In the Individual condition, we stated that the higher-scoring applicant "displays ability and potential" and that the lower-scoring applicant "displays ability, potential, and faced adversity." We also described their differences: the higher-scoring applicant "has stronger scores on attributes that signal ability and potential (e.g., SAT scores, GPA) than [the other applicant]…" and the lower-scoring applicant "has weaker scores on attributes that signal ability and potential (e.g., SAT scores, GPA) and has faced more adversity than [the other applicant]."

The Policy condition had the same format, but with slightly different language. We described the policy that would favor higher-scoring type of applicants as considering "ability and potential" and the policy that would favor the lower-scoring type of applicants as considering "ability, potential, and adversity that applicants face." For the policy that favored higher-scoring applicants we specified that "applicants admitted under [this policy] generally have stronger scores on attributes that signal ability and potential (e.g., SAT scores, GPA) than applicants admitted under [the other policy]," whereas for the policy that favored lower-scoring applicants we specified that "applicants admitted under [this policy] generally have weaker

scores on attributes that signal ability and potential (e.g., SAT scores, GPA) and have faced more adversity than applicants admitted under [the other policy]." After examining the table of information, participants selected one of the two applicants or policies (forced choice: A, B) and indicated which option was the fairer one to select (forced choice: A, B).

Participants also indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with two statements pertaining to moralization and two statements pertaining to relying on objective information. These items were identical to those from Study S3. We created composites of the moralization (label: *Moralization composite*; $r = .34$) and objective information (label: *Objective information composite*; $r = .60$).

Participants also answered exploratory questions about their attributions of each applicant. They reported how much of each applicant's "level of achievement is the result of their…": (a) "hard work," (b) "abilities," (c) "motivation," (d) "household income," (e) "academic opportunities," and (f) "luck" (1 = *None at all* to 5 = *A great deal*). The labels for these items are identical to the item (e.g., the label for the hard work item is *Hard work*). Participants also reported each applicant's level of achievement (label: *Achievement*; 1 = *One of the worst* to 5 = *One of the best*).

At the end of the study, participants provided open-ended answers for why they chose the applicant or policy that they did and what the difference between the two options was. They also provided demographic information (age, gender, political attitudes, and political affiliation).

Results

Choice. We conducted a logistic regression on the choice of applicant or policy (1 = higher-scoring, higher-income option; 0 = lower-scoring, lower-income option), using Decision Type (1 = Individual, 0 = Policy), Adversity (1 = Salient, 0 = Not Salient), and the Decision

Type x Adversity interactions as independent variables. The coefficient on Decision Type was significant (b = 1.116, p < .001), suggesting that participants in the Not Salient condition who made a choice about individuals were significantly more likely to select the higher-scoring, higher-income option (73.83%) than participants who made a choice about policies (48.03%). The coefficient on Adversity was not significant (b = -0.153, p = .509). In other words, the choices of participants who made a policy decision in the Not Salient (48.03%) and Salient (44.22%) conditions were not significantly different. The absence of an adversity effect on the policy choice suggests that policies may naturally call adversity to mind. Importantly, we found a significant interaction of Decision Type x Adversity (b = -1.033, p = .002), indicating that making adversity salient shifted choices significantly more in the Individual condition (Not Salient: 73.83% vs. Salient: 46.26%) than in the Policy condition. Making adversity salient aligned choices in the Individual condition with choices in the Policy condition. Indeed, choices about individuals when adversity was salient were not significantly different from choices about policies when adversity was salient (χ^2 [1, 294] = 0.13, p = .725), nor from choices about policies when adversity was not salient (χ^2 [1, 299] = 0.09, p = .760).

Fair choice. We conducted a logistic regression on the choice of the fair applicant or policy to select (1 = higher-scoring, higher-income option; 0 = lower-scoring, lower-income option), using the same independent variables as the regression on Choice, and found the same pattern of results. The coefficient on Decision Type was significant (b = 1.480, p < .001), suggesting that participants who made a choice about individuals in the Not Salient condition were significantly more likely to select the higher-scoring, higher-income option as the fair choice (75.17%) than participants who made a choice about policies (40.79%). The coefficient on Adversity was not significant (b = 0.250, p = .284). In other words, the fair choice judgments

of participants who made a policy decision in the Not Salient (40.79%) and Salient (46.94%) conditions were not significantly different. As with Choice, we found a significant interaction of Decision Type x Adversity (b = -0.986, p = .004), indicating that making adversity salient in the Individual condition shifted fair choice judgments significantly more in the Individual condition (Not Salient: 75.17% vs. Salient: 59.18) than in the Policy condition.

Moralization. We conducted an ANOVA on the moralization composite, with factors for Decision Type, Adversity, and their interaction. We found main effects of Decision Type (p = .003) and of Adversity (p < .001), but not a significant interaction (p = .292). When Adversity was not salient, participants who made an individual decision did not agree more or less with the moralization measure than participants who made a policy decision (p = .001). However, despite the lack of a significant overall ANOVA, when Adversity was salient, participants who made an individual decision agreed less with the moralization measure than participants who made a policy decision (p = .001).

Objective Information. We conducted an ANOVA on the objective information composite, with factors for Decision Type, Adversity, and their interaction. We found main effects of Decision Type (p = .002) and of Adversity (p < .001), as well as their interaction (p = .021). When Adversity was not salient, participants who made an individual decision agreed more with objective information measure than participants who made a policy decision (p < .001). However, when Adversity was salient, there was no significant difference between participants who made an individual versus policy decision (p = .618).

SM 2.11. STUDY S11 (RELIANCE ON VALUES OR OBJECTIVE INFORMATION)

In Study S11, we sought to test whether the policy frame, relative to the individual frame, causes decision-makers to favor lower-scoring, disadvantaged options by leading them to rely on their values more than on objective information. To do so, we directly manipulated participants' reliance on specific values or objective information. We preregistered this study on AsPredicted.org (https://aspredicted.org/ZRZ_KFV).

Methods

Participants. We recruited participants on MTurk for a small monetary sum. As indicated in our preregistration, we excluded participants who either twice failed an attention check before being assigned to a condition, did not pass a display compatibility check, or completed the survey in under 120 seconds. We ended with a final sample of 595 participants (M_{age} = 38.59, *SD* = 10.90; 48.57% female).

Procedure and materials. We randomly assigned participants to a condition from a 2 (Decision Type: Individual, Policy) x 3 (Reliance: Control, Objective Information, Values) between-subjects design. As in Study 1, we asked all participants to imagine that they were the head of admissions at a four-year university and had to make decisions about which applicants to admit. In all conditions, we presented participants with a similar table of information as in Study 1 (this table omitted information about GPA and extracurriculars), and additionally we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option was labeled "A" and appeared on the left.

The Control versions of the Individual and Policy conditions were identical to those we tested in Study 1 (except for the within-subjects portion); hence, Control participants simply selected an option from the two applicants or policies presented. In the Objective Information and Values conditions, participants read an additional instruction before seeing the information

about the applicants or policies. Specifically, participants read, "When making admissions decisions, people may feel conflicted between relying on the objective achievements of the applicants or on their values around equality and diversity." This statement was followed by a prompt tailored to their condition: Objective Information participants were prompted to choose "only by relying on the objective achievements of the applicants presented," whereas Values participants were prompted to choose "only by relying on values around equality and diversity". We chose to focus participants in the Values condition specifically on the values of equality and diversity, rather than on values more generally because exploratory measures we collected in previous studies suggested these were the values motivating people's choices in the context of college admissions.[8]

After reading the prompts, participants saw the table of information describing the applicants or policies. Then, participants chose one of the two options (forced choice: A, B). Participants in the Objective Information and Values condition were reminded to make their choice by relying on the "objective achievements of the applicants presented" or "on values around equality and diversity." All participants then indicated how likely they were to pick either option (1 = *Much more likely to choose A,* 4 = *Equally Likely*, 7 = *Much more likely to choose B*) and, "regardless of what [they] chose before," which option was fairer to choose (forced choice: A, B).

Values and objective information. Next, we measured the extent to which participants made the choice based on their values or objective information. They first indicated whether their choice was based more on their values or on objective information (1 = *Much more my values,* 4 = *Equally,* 7 = *Much more objective information*). Then, they indicated their agreement

[8] In a pilot study we asked participants to make their decision "only by relying on the values that [they] believe are important for society," and we found that this more general prompt did not influence choice.

(1 = *Strongly Disagree*, 4 = *Neither Agree Nor Disagree*, 7 = *Strongly Agree*) with six statements: "My choice was based on…" (1) "…objective criteria.", (2) "…concrete and factual reasons.", (3) "…my value of increasing diversity.", (4) "…my value of giving opportunities to those in need.", (5) "…my value of creating equality.", and (6) "…the values I think are important for society."[9] Consistent with the preregistration, we combined these seven items into a *Values-versus-objective-information* composite ($\alpha = .89$). A higher value on the composite indicates that participants relied more on objective information, rather than their values, when making their choice. For completeness, we report analyses for composites of the values items (#3-6, label: *Values composite*; $\alpha = .90$) and of the objective information items (#1-2; label: *Objective composite*; $r = .70$) in SM 3. At the end of the study, participants provided demographic information (age, gender, and political attitudes) and were paid.

Results and discussion

Choice. When not prompted to rely on any information (Control), 67.02% of participants making a choice between individuals chose the higher-scoring, higher income option, whereas only 40.21% of participants making a choice between policies selected that option (χ^2 [1, 191] = 13.79, $p < .001$). When prompted to rely on objective information, 82.47% of participants making a choice between individuals chose the higher-scoring, higher income option, whereas only 69.70% of participants making a choice between policies selected that option (χ^2 [1, 196] = 4.39, $p = .036$). When prompted to rely on their values, 47.62% of participants making a choice between individuals chose the higher-scoring, higher income option, whereas only 29.13% of participants making a choice between policies selected that option (χ^2 [1, 208] = 7.51, $p = .006$).

[9] We conceptualized the values items as reflecting specific values prioritized when making these decisions. Some of these items (3 and 5) are similar in nature to the macrojustice items of Studies 4, S5, and S7. However, because the items in Study S11 were asked differently than in those studies and encompass a wider set of values (i.e., need and what is important for society), we did not combine them into a Macrojustice composite.

Choice regression. We conducted a logistic regression on choice (1 = higher-scoring, higher-income option; 0 = lower-scoring, lower-income option), using Decision Type (1 = Individual, 0 = Policy), Objective Information (1 = Reliance on Objective Info condition, 0 = all other), Values (1 = Reliance on Values condition, 0 = all other), Decision Type x Objective Information, and Decision Type x Values as independent variables. We present the regression results in Table S3, along with the results of additional regressions predicting likelihood, fair choice, and the values-versus-objective information composite.

Likelihood regression. We conducted an OLS regression of the likelihood to choose the higher-scoring, higher-income option over the lower-scoring, lower-income option, with the same factors used in the logistic regression of choice. The pattern of results, which we present in Table S3, is similar to the Choice regression.

Fair choice regression. We conducted a logistic regression on fair choice, with the same factors used in the logistic regression of choice. We present the regression results in Table S3. The overall pattern of results indicates that asking participants to rely on their values versus objective information did not shift what participants thought was fair to select. That said, participants' views about the fair option to select were influenced by whether they were making an individual or policy decision (consistent with the rest of the studies).

Values-versus-objective-information regression. We conducted an OLS regression on the Values-versus-Objective-Information composite, using the same independent variables we used in the logistic regression of choice. The interactions between the type of decision and the reliance prompts were not significant. However, we find directional support for the predicted pattern of results. Without prompting, participants naturally relied more on objective information than the values of diversity and equality when making decisions about individuals than when making

policy decisions. When we prompted participants to rely on objective information, this did not change the choices of participants making a decision about individuals (because relying on this information is relatively consistent with the way people make individual decisions), but it did change the choices of participants making decision about policies. More participants selected policies that favored higher-scoring applicants, thereby somewhat attenuating the individual-versus-policy divergence. Conversely, when we prompted participants to rely on the values of equality and diversity, this did not change the decisions of participants making policy decisions (because relying on these values is relatively consistent with the way people make policy decisions), but it did change the decisions of participants making decisions about individuals. More participants selected the specific lower-scoring individual, thereby somewhat attenuating the individual-versus-policy divergence.

TABLE S8

Results of Regressions Predicting Choice, Likelihood, Fair Choice, and Reliance on Values versus Objective Information (Study S11)

	Model 1: Logistic DV = Choice				Model 2: OLS DV = Likelihood				Model 3: Logistic DV = Fair Choice				Model 4: OLS DV = Values-versus-Objective			
Variable	***b***	***SE***	***p***	***95% CI***	***b***	***SE***	***p***	***95% CI***	***b***	***SE***	***p***	***95% CI***	***b***	***SE***	***p***	***95% CI***
Decision Type	1.106	0.302	< .001	[0.515, 1.697]	0.842	0.286	.003	[0.281, 1.404]	0.828	0.302	.006	[0.236, 1.419]	0.796	0.198	< .001	[0.407, 1.186]
Objective	1.230	0.301	< .001	[0.640, 1.820]	1.167	0.282	< .001	[0.613, 1.722]	-0.081	0.286	.778	[-0.641, 0.480]	1.033	0.196	< .001	[0.648, 1.417]
Values	-0.492	0.300	.101	[-1.080, 0.095]	-0.430	0.280	.124	[-0.980, 0.119]	-0.077	0.283	.787	[-0.632, 0.478]	0.172	0.194	.375	[-0.209, 0.553]
Individual x Objective Information	-0.390	0.458	.395	[-1.289, 0.508]	-0.329	0.402	.413	[-1.118, 0.460]	0.077	0.424	.856	[-0.754, 0.908]	-0.745	0.279	.008	[-1.292, -0.198]
Individual x Values	-0.312	0.420	.457	[-1.135, 0.511]	-0.229	0.396	.564	[-1.007, 0.549]	-0.121	0.415	.769	[-0.934, 0.690]	-0.916	0.275	.001	[-1.455, -0.376]
Constant	-0.397	0.207		[-0.803, 0.009]	3.722	0.201		[3.328, 4.116]	-0.020	0.203		[-0.419, 0.377]	3.236	0.139		[2.962, 3.509]
(Pseudo*) R Squared	0.11*				0.12				0.03*				0.10			

Note. For each model, the independent variables are Decision Type (1 = Individual, 0 = Policy), Objective Information (1 = Reliance on Objective Information condition, 0 = all other), Values (1 = Reliance on Values condition, 0 = all other), and the interactions of Decision Type x Objective Information and Decision Type x Values. Models 1-3: Higher values correspond to a greater preference for the higher-scoring, higher-income option. Model 4: Higher values indicate greater reliance on objective information than on values. 95% confidence intervals refer to unstandardized betas.

SM 2.12. STUDY S12 (MECHANISM PILOT 1)

We measured several factors that may explain the divergence between individual and policy decisions. Specifically, we measured people's focus on the decision-making process versus the outcome of the decision, the extent to which people saw the choice as moral and as objective, and how consequential they perceived the choice to be. We preregistered this study on AsPredicted.org (https://aspredicted.org/ADJ_XKO).

Methods

Participants. We recruited participants on Amazon Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check before being assigned to a condition or who completed the survey in under 90 seconds. We ended with a final sample of 298 participants (M_{age} = 37.90, SD = 12.55; 47.32% female).

Procedure and materials. Participants were randomly assigned to make either an Individual or Policy decision. As in Study 1, participants making an Individual decision selected one of two individual applicants to admit, whereas participants making a Policy decision selected one of two admissions policies. We also used the same stimuli as in Study 1. Specifically, we used the same table of information to describe the applicants (in the Individual condition) and types of applicants (in the Policy condition). Additionally, we counterbalanced whether the higher-scoring, higher-income option or the lower-scoring, lower-income option would be labeled "A" and appear on the left.

After reading the information about the options (applicants or policies) and examining the table of information, all participants indicated which of the two options they would select and which option would be the fairer one to select (forced choice: A, B). Participants then explained, in a free-response format, why they chose the particular applicant or policy that they did.

Next, participants answered questions about the first mechanism: their focus on the decision-making process versus the outcome of the decision. We explained to participants that, "When making a decision, you can focus on the decision-making process (e.g., what criteria you use, how you compare options) or the decision outcome (e.g., the final option you choose, the implications of that choice)." Participants then answered three questions, which were presented in a random order: "For your choice of [applicant/policy], to what extent were you focused on using a fair decision process?" (label: *Process focus*; 1 = *Not at all,* 7 = *Extremely*), "For your …, to what extent were you focused on arriving at a fair outcome" (label: *Outcome focus*; 1 = *Not at all,* 7 = *Extremely*), and "For your …, how important was the decision making-process versus the decision outcome?" (label: *Proc. vs. outcome*; 1 = *Process Extremely More Important,* 4 = *Process and Outcome Equally Important*, 7 = *Outcome Extremely More Important*).

Participants then answered questions about the second potential mechanism: the extent to which people saw the choice as moral and as objective. Participants indicated their agreement (1 = *Strongly disagree* to 7 = *Strongly agree*) with four moralization and two objective information items, all presented in random order. Two of the moralization items and the two objective information items were identical to those from Studies S3 and S10. The two additional moralization items were: "There are very important ethical aspects to this decision." and "This decision clearly does not involve ethics or moral issues" (reverse-coded). We created composite measures of the moralization (label: *Moralization composite;* $\alpha = .79$) and objective information items (label: *Objective information composite;* $r = .59$). We also created a values-vs-objective-information composite of these items (label: *Values-versus-Objective Information*), as in the main manuscript ($\alpha = .74$).

Next, participants answered questions about the third potential mechanism, how consequential they perceived the choice to be. Participants answered three questions, which were presented in random order: “How important is this decision?” (label: *Importance;* 1 = *Not at all*, 7 = *Extremely*), “How consequential is this decision?” (label: *Consequential;* 1 = *Not at all*, 7 = *Extremely*), and “How many people will be affected by this decision?” (label: *Scope;* 1 = *No people at all, 4 = A moderate amount of people,* 7 = *A large amount of people*). Finally, participants provided demographic information (age, gender, ethnicity, political attitudes, highest level of schooling completed, and their household’s current income level). SM 3 contains all results for non-choice and non-fairness measures for this study.

SM 2.13. STUDY S13 (MECHANISM PILOT 2)

In Study S12 we did not find a significant difference in the choices made about individuals and policies. Hence, we were not able to test whether the proposed factors underlie the divergence between individual and policy decisions. In order to test the proposed factors, in Study S13 we replicated the procedure of Study S12 with some slight modifications. We preregistered this study on AsPredicted.org (https://aspredicted.org/TNZ_AZG).

Methods

Participants. We recruited participants on Amazon Mechanical Turk for a small monetary sum. As indicated in our preregistration, we excluded participants who twice failed an attention check or who completed the survey in under 90 seconds. We ended with a final sample of 287 participants (M_{age} = 34.73, *SD* = 11.29; 46.69% female).

Procedure and materials. Study S13's procedure and materials are identical to Study S12's with the modifications described ahead. We revised the measures for the proposed factors. We changed the order in which we measured the proposed factors to first measure the extent to which people saw the choice as moral and as objective. We dropped two of the moralization statements that required reverse coding ("Morality is irrelevant to my choice of [applicant/policy]." and "This decision clearly does not involve ethics or moral issues.") and we added one new item ("My choice of [applicant/policy] is connected to my beliefs about fundamental right and wrong."). These moralization items we dropped are the ones we used in Studies S3 and S10. We created composite measures of the two remaining moralization items (label: *Moralization composite;* α = .83) and objective information items (label: *Objective information composite; r* = .57). We also created a values-vs-objective-information composite of these items (label: *Values-versus-Objective Information;* α = .45).

We changed the questions that measured participants' focus on the decision-making process versus the outcome of the decision. Instead of asking three questions with an introductory statement, we simply asked two questions: "When making your choice, to what extent were you focused on your decision process (for example, how to compare the information across options)?" (label: *Process focus*) and "When making your choice, to what extent were you focused on the outcome of your choice?" (label: *Outcome focus*). Participants answered these questions using 7-point Likert scale (1 = *Not at all,* 4 = *Moderately,* 7 = *Extremely*). We repeated the same questions as in Study S12 about the consequence (label: *Consequential*), importance (label: *Importance*), and scope (label: *Scope*) of the decision, except that the questions now referred to "your decision" instead of "the decision."

We also measured participants' agreement with the Social Justice Moral Motives (Janoff-Bulman & Carnes, 2016). Participants indicated their agreement (1 = *Strongly disagree,* 4 = *Neither agree, nor disagree,* 7 = *Strongly agree*) with the following five statements: "It is our responsibility, not just a matter of personal preference, to provide for groups worse off in society", "Giving to groups worse off in society does not make those groups too dependent on help", "It is important for those who are better off to help provide resources for the most vulnerable members of society", "Increased economic equality is ultimately beneficial to everyone in society", and "In the healthiest societies, those at the top should feel responsible for improving the well-being of those at the bottom." We created a composite measure (α = .91).

Participants provided the same demographic information as before (age, gender, ethnicity, political attitudes, highest level of schooling completed, and their household's current income level), and additionally indicated their political attitudes for economic issues and for social issues. SM 3 contains all results for non-choice and non-fairness measures for this study.

SM 2.14. STUDY S14 (SURVEY OF MANAGEMENT FACULTY)

We surveyed management faculty about their previous experience making hiring decisions. We were interested in whether they had previously faced a tradeoff between achievement and circumstances when hiring faculty, whether their decision when facing said tradeoff was guided by macrojustice and microjustice concerns, and who they recalled choosing when facing said tradeoff. Because most faculty likely did not have experience making policy decisions (an earlier pilot suggested this was the case), we only asked them to focus on experiences making individual selection decisions.

Methods

Participants. We emailed 500 faculty from top management department in the U.S. (per the UTD Top 100 Business School Research Rankings). Our email explained that were interested in surveying them about "faculty hiring." We collected data from 114 participants, 8 of which (7%) reported not having been involved or having a conversation about hiring faculty to their department, which prevented them from completing the rest of the survey. Of the remaining 106 participants, the vast majority reported recent involvement in hiring (90% were involved within the last three years). We asked them recall a time they faced a tradeoff between "candidates with higher objective metrics of success (e.g., many publications, excellent methodological training) or candidates with marginally lower metrics of success who may have been disadvantaged in some way (e.g., they lacked research funding, had a limited stipend, belong to a group or social category that faces more obstacles in academia)". Eighty-six percent of the management faculty could recall making faculty hiring decisions that involved this tradeoff, which is the same one we represent in our studies. We purposely did not collect demographic information.

Procedure and materials. We describe to participants the sort of tradeoff between achievement and circumstances that we were interested in. Specifically, participants read the

following text, before responding whether they had ever faced a tradeoff like it when making hiring decisions (Yes, No, Other).

> "When making final hiring decisions, departments are often faced with a small pool of highly qualified candidates, several of whom meet the criteria for hiring. A department may only be able to hire one of the candidates, and choosing among them can be particularly difficult because of differences in their achievements and circumstances.
>
> Hiring committees may sometimes need to trade off choosing candidates with higher objective metrics of success (e.g., many publications, excellent methodological training) or candidates with marginally lower metrics of success who may have been disadvantaged in some way (e.g., they lacked research funding, had a limited stipend, belong to a group or social category that faces more obstacles in academia).
>
> This sort of tradeoff is, of course, only one part of a broader set of factors that influence how hiring and other selection decisions are made. Nevertheless, it can be particularly relevant and pronounced in some situations. For example, in college admissions, socioeconomic income is strongly correlated with SAT scores (Freedle, 2003) and measures of essay quality (Alvero et al., 2021), such that officers sometimes face a tradeoff between selecting applicants with the highest scores who tend to be more socioeconomically advantaged and applicants with lower scores who tend to be socioeconomically disadvantaged."

Depending on participants' response, we asked them questions about a past experience (if they indicated Yes) or a hypothetical one (if they indicated No or Other). Participants who were asked about a past experience were asked whether, when they faced a tradeoff between candidates' achievements and their circumstances, they hired a candidate with exceptional metrics of success who was relatively advantaged or a candidate with strong metrics who was relatively disadvantaged (or selected the "Other" option). Participants who were asked about a hypothetical experience were asked the same question but framed as who they would choose *if* they faced such a tradeoff.

We counterbalanced whether participants first indicated what choice they made (or would make) or first answered questions about whether their choice was (or would be) guided by macrojustice and microjustice considerations. In other words, half of participants answered the question about who they hired (or would hire) before answering questions about whether microjustice and macrojustice considerations guided (or would guide) their choice, while the

other half of participants first answered questions about the considerations and then about the choice.

The microjustice and macrojustice items were similar to those in Study 3. We asked participants how much their preferred course of action was guided four macrojustice considerations and four microjustice considerations (1 = *Not at all motivated by it*, 4 = *Moderately motivated by it*, 7 = *Extremely motivated by it*). The four macrojustice items were: "Making sure that there is a diverse population of faculty in my department or field," "Creating greater equality in the jobs disadvantaged and advantaged candidates receive," "Balancing the number of jobs given to disadvantaged and advantaged candidates," and "Making sure that there are not too few disadvantaged faculty in my department or field." The four microjustice items were: "Selecting candidates based on their hard work," "Selecting candidates based on their qualifications," "Selecting applicants based on their potential to succeed as faculty," and "Selecting candidates based on their accomplishments." We combined the four macrojustice items into a composite (α = .80) and the four microjustice items into a composite (α = .71).

At the end of the survey participants indicated when they were last involved in the hiring process at their department and their greatest involvement in their department's hiring process.

Results and discussion

Tradeoff. The vast majority of participants (84.91%) reported facing a tradeoff between achievement and applicants' circumstances when making faculty hiring decisions. Only 13.21% reported not facing such a tradeoff and 1.89% selected the Other option. We focus our analyses only on participants who reported facing the tradeoff.

Choice. The majority of participants reported selecting the candidate with lower metrics of success who was disadvantaged (58.89%), whereas only a quarter reported selecting the

candidate with higher metrics who was advantaged (26.67%). The remaining participants (14.44%) selected the Other option.

Macrojustice and microjustice. On average, participants reported being more than just moderately motivated (the midpoint of the scale) by both macrojustice (M = 4.40, SD = 1.46) and microjustice concerns (M = 5.69, SD = 0.81). They reported being more motivated by microjustice than macrojustice ($p < .001$).

We also examined whether participants' reported choice of candidate was predicted by their macrojustice and microjustice motivations. For this analysis we omitted participants who selected the Other option. We conducted an OLS regression predicting choice (1 = higher-metrics, advantaged candidate; 0 = lower-metrics, disadvantaged candidate) with factors for the macrojustice and microjustice composites. We found that participants macrojustice concerns significantly predicted selecting the lower-metrics, disadvantaged candidate (b = -0.12, $p < .001$), but their microjustice concerns did not (b = -0.04, p = .502). This was also the case when we conducted separate regressions with only macrojustice or microjustice as predictors.

These results provide evidence that decision makers who make hiring decisions on a regular basis, and the majority of whom have done so recently, report relying on both macrojustice and microjustice concerns. Moreover, those concerns predict who they report selecting when they faced a tradeoff between candidates' achievements and their circumstances, which is the same tradeoff that we present across our studies.

SM 2.15. STUDY S15 (SURVEY OF ADMISSIONS OFFICERS)

We surveyed admissions officers about their experience with college admissions. We were primarily interested in learning how often they faced a tradeoff like the one presented in our studies, as well as in their explanations for how they went about making decisions that presented this tradeoff and how they thought about the fair thing to do in those situations. Several of the questions were qualitative, which we describe ahead but do not analyze; we only analyze the quantitative measure. Additionally, we collected measures that are relevant for another research project; we describe those measures here but do not report their results. Nevertheless, we make all of the data available for readers interested in analyzing it further.

Methods and results

Participants. We recruited admissions officers by sharing a recruitment blurb with admissions officers we knew. These officers shared those recruitment blurbs with fellow admissions officers, in some cases via email and in other cases via private Facebook groups (including a "women in admissions" group). We incentivized participating in our survey via a $10 Amazon gift card. We described the survey as relevant to "admissions professionals (anyone who currently works or has previously worked in college admissions)" and as asking about "your experience making admissions decisions." We had 35 participants complete our survey in full (M_{age} = 35.03, SD = 7.66; 85.71% female; 77.14% White; 100% slightly to extremely liberal). All participants reported at least 2 years of work experience in an admissions department (M = 10.34, SD = 6.24).

Procedure, materials, and results. Participants answered multiple questions about their experience with college admissions. First we asked participants whether, when making a decision about who to move forward in the admissions process, they had ever had to choose "between (a) an applicant with high qualifications who came from a relatively privileged

background versus (b) an applicant with lower qualifications from a disadvantaged background (e.g., socioeconomically, racially)?" (Yes, No, Other). The majority of participants indicated Yes (51.45%), with the remaining similarly split between No (25.71%) and Other (22.86%).

Participants indicated how often they faced decisions like the one above (1 = *Never*, 3 = *Sometimes*¸5 = *All the time*). Among participants who responded Yes to the earlier question, the average response was greater than sometimes (M = 3.56, SD = 0.78).

Participants then explained (in text) what is the fairest way to make a decision like the one we first asked about, as well as whether there are other perspectives on the fairest way to make the decision. We do not analyze these qualitative responses.

We asked participants to "think of a recent example where you personally had to choose between (a) an applicant with high qualifications who came from a relatively privileged background versus (b) an applicant with lower qualifications from a disadvantaged background." Then we asked participants who they moved forward in the admissions process in the specific case they had in mind. Among participants who reported having faced a decision like this before (based on the first question), half (50%) reported moving the lower-qualified, disadvantaged applicant forward, a third reported moving the higher-qualified, advantaged applicant forward (33.33%), and the remaining (16.67%) selected Other. Then participants explained (in text) how they decided which applicant to move forward.

Participants then reported whether the decision was difficult to make (1 = *Not at all difficult*, 4 = *Difficult*, 7 = *Extremely difficult*). Among participants who reported having faced a decision like this before (based on the first question), on average they reported that the decision was difficult (M = 4.11, SD = 1.68). Then participants explained (in text) why the decision was difficult or not.

Next, participants reported whether their choice of applicant to move forward was motivated by a desire to… "create equality", "increase diversity", and "recognize applicants for their accomplishments" (1 = *Not motivated by it all, 4 = Motivated by it,* 7 = *Extremely motivated by it*). These are the same items that we presented to participants in Study 4. Among participants who reported having faced a decision like this before (based on the first question), they reported being motivated by all of these (equality: $M = 3.89$, $SD = 2.00$; diversity: $M = 4.11$, $SD = 1.81$; equality: $M = 4.00$, $SD = 1.68$).

Finally, participants answered questions pertaining to another research project. Participants imagined looking through an application and then learning that the applicant was advantaged or disadvantaged. Participants reported whether this information would change their impressions of the applicant, and they explained why. Participants repeated this same process but now imagining that the applicant was disadvantaged or advantaged (whichever option they had not seen). After this, participants reported demographic information and provided details about their work experience and institution.

These results provide evidence that college admissions officers face the tradeoffs that we present in this research project. Moreover, they also report that macrojustice and microjustice are relevant considerations to their decisions.

SM 3. RESULTS OF ALL NON-CHOICE MEASURES COMPARING POLICY AND INDIVIDUAL DECISIONS

Category	Label	Study	*r* / *α*	Conditions	Individual		Policy		*t*	*p*	
					M	*SD*	*M*	*SD*			
Reliance on Justice Principles	Macrojustice	3	.93	Standard	3.89	1.93	4.14	1.91	1.80	.072	†
		4	.82	Standard	3.75	1.84	4.56	1.76	5.99	<.001	***
				Macrojustice	4.53	1.92	4.95	1.79	2.85	.005	**
		S5	.81	Policy-Narrow	3.57	1.94	3.80	1.93	1.47	.142	
				Policy-Abstract			4.00	2.00	2.65	.008	**
		S7	.76	Standard	3.59	1.99	4.34	1.90	3.22	.001	**
		S8	.84	Standard	3.74	2.15	4.37	1.98	2.32	.021	*
	Microjustice	3	.87	Standard	5.79	0.97	5.51	1.22	3.63	<.001	***
		4		Standard	5.30	1.53	4.95	1.46	3.15	.002	**
				Macrojustice	4.57	1.70	4.60	1.61	0.20	.845	
		S7		Standard	5.26	1.52	5.16	1.51	0.56	.579	
		S8		Standard	5.51	1.39	5.35	1.61	0.86	.390	
	Macro concerns	S8	.83	Standard	4.79	1.47	5.00	1.22	1.21	.226	
	Micro concerns	S8	.83	Standard	4.70	1.32	4.86	1.20	0.94	.346	
	Procedural	S8	.85	Standard	5.14	1.10	5.24	0.98	0.72	.475	
	Distributive	S8	.85	Standard	5.27	1.01	5.39	1.13	0.88	.378	
Reliance on Values and Objective Information	Values-versus-objective-information	S6		Standard	5.12	1.61	4.22	1.64	5.43	<.001	***
				Rule			4.19	1.63	5.79	<.001	***
		S11	.89	Standard	4.03	1.30	3.24	1.17	4.45	<.001	***
	Values	1[a]		Standard	0.07	2.91	--	--	0.33	.743	
		S11	.90	Standard	4.03	1.50	3.13	1.44	4.21	<.001	***
	Objective Information	1[a]		Standard	1.05	2.48	--	--	5.76	<.001	***
		S3	.58	Policy-Specific	5.71	1.12	5.20	1.20	4.46	<.001	***
				Policy-Range			5.22	1.16	4.35	<.001	***
		S4	.57	Policy-Narrow	5.57	1.23	5.37	1.17	1.65	.101	
				Policy-Broad			5.50	1.03	0.58	.562	
				Policy-Abstract			5.27	1.10	2.56	.011	*

(continued; 2 of 5)

Category	Label	Study	*r* / *α*	Conditions	Individual		Policy		*t*	*p*	
					M	*SD*	*M*	*SD*			
Reliance on Values and Objective Information (continued)	Objective Information (continued)	S5	.62	Policy-Narrow	5.94	1.16	5.66	1.21	2.77	.006	**
				Policy-Abstract			5.56	1.27	3.70	<.001	***
		S6	.59	Standard	5.50	1.21	5.26	1.11	2.01	.045	*
				Rule			5.17	1.14	2.79	.005	**
		S10	.60	Not Salient	5.80	1.06	5.28	1.18	3.99	<.001	***
				Salient	5.20	1.23	5.13	1.22	0.50	.618	
		S11	.70	Standard	5.57	1.23	4.99	1.40	3.04	.003	**
		S12	.59	Standard	5.53	1.24	5.25	1.16	1.97	.049	*
		S13	.57	Standard	5.62	1.17	5.39	1.00	1.75	.081	†
Moralization of Choice	Moralization	2	.79	Standard	4.68	1.38	4.87	1.35	1.35	.178	
		S3	.40	Policy-Specific	3.89	1.35	4.99	1.33	8.16	<.001	***
				Policy-Range			4.93	1.24	8.11	<.001	***
		S4	.72	Policy-Narrow	4.96	1.29	5.15	1.23	1.52	.129	
				Policy-Broad			5.37	1.13	3.36	<.001	***
				Policy-Abstract			5.43	1.10	3.97	<.001	***
		S6	.83	Standard	4.39	1.30	5.11	1.27	5.41	<.001	***
				Rule			5.05	1.23	5.27	<.001	***
		S10	.34	Not Salient	4.19	1.25	4.36	1.49	1.06	.292	
				Salient	4.44	1.24	4.90	1.18	3.25	.001	**
		S12	.79	Standard	4.36	1.28	4.52	1.39	1.07	.287	
		S13	.83	Standard	4.93	1.31	5.00	1.33	0.41	.682	
Focus of Choice	Process-vs-Outcome	S12		Standard	4.36	1.56	4.42	1.74	0.30	.762	
	Process	S12		Standard	2.42	1.26	2.80	1.51	2.38	.018	*
		S13		Standard	5.39	1.51	5.30	1.36	0.51	.608	
	Outcome	S12		Standard	5.47	1.30	5.22	1.51	1.52	.130	
		S13		Standard	5.15	1.63	5.42	1.38	1.51	.133	
	Justification	1[a]		Standard	0.51	2.57	--	--	2.69	<.001	***
	Adversity Faced	1[a]		Standard	0.67	2.71	--	--	3.37	.001	**
	Identity	1[a]		Standard	0.42	0.19	--	--	2.18	.031	*

(continued; 3 of 5)

Category	Label	Study	*r* / *α*	Conditions	Individual		Policy		*t*	*p*	
					M	*SD*	*M*	*SD*			
Perceptions of Choice	Zero-sum	S3		Policy-Specific	4.02	1.22	4.01	1.05	0.08	.938	
				Policy-Range			4.10	1.08	0.74	.462	
	Guide	S5		Policy-Narrow	4.81	1.52	5.66	1.15	7.56	<.001	***
				Policy-Broad			5.83	1.01	9.48	<.001	***
	Statement of Self	S5		Policy-Narrow	3.57	1.82	4.07	1.69	3.45	<.001	***
				Policy-Abstract			4.10	1.78	3.57	<.001	***
	Statement of Company	S5		Policy-Narrow	4.57	1.79	5.19	1.48	4.54	<.001	**
				Policy-Abstract			5.17	1.53	4.38	<.001	***
	Statement	S6		Standard	5.02	1.32	5.45	1.20	3.36	<.001	***
				Rule			5.40	1.26	3.00	.003	**
	Consequential	S4		Policy-Narrow	4.91	1.43	4.80	1.58	0.73	.465	
				Policy-Broad			5.01	1.53	0.69	.488	
				Policy-Abstract			5.17	1.39	1.90	.059	†
		S12		Standard	4.90	1.51	5.09	1.35	1.10	.274	
		S13		Standard	5.01	1.51	5.07	1.36	0.35	.727	
	Importance	S4		Policy-Narrow	5.44	1.24	5.31	1.38	1.02	.306	
				Policy-Broad			5.55	1.29	0.87	.387	
				Policy-Abstract			5.70	1.05	2.23	.026	*
		S12		Standard	5.48	1.15	5.51	1.08	0.26	.794	
		S13		Standard	5.41	1.24	5.43	1.21	0.12	.905	
	Scope	S4		Policy-Narrow	3.35	1.49	3.29	1.67	0.42	.672	
				Policy-Broad			5.21	1.64	11.97	<.001	***
				Policy-Abstract			5.49	1.35	15.16	<.001	***
		S12		Standard	3.23	1.31	4.72	1.63	8.69	<.001	***
		S13		Standard	3.25	1.39	4.68	1.71	7.84	<.001	***
Perceptions of Applicant	Achievement: Advantaged[b]	S10		Not Salient	3.91	0.04	3.79	0.50	2.03	.044	*
				Salient	3.75	0.58	3.73	0.49	0.36	.718	
	Achievement: Disadvantaged[b]	S10		Not Salient	1.02	0.05	4.58	0.56	1.25	.211	
				Salient	4.43	0.56	4.41	0.56	0.28	.777	

(continued; 4 of 5)

Category	Label	Study	*r* / *α*	Conditions	Individual		Policy		*t*	*p*	
					M	*SD*	*M*	*SD*			
Perceptions of Applicant (continued)	Qualified	S3		Policy-Specific	5.47	1.03	5.37	1.23	0.87	.383	
				Policy-Range			5.19	1.31	2.38	.018	*
Causes of Applicant's Achievement	Skill/Effort	2		Standard	4.74	1.92	5.34	1.65	3.19	.002	**
	Effort	S3		Policy-Specific	4.21	1.28	3.82	1.47	2.83	.005	**
				Policy-Range			3.72	1.57	3.48	<.001	***
	Hard work: Advantaged[b]	S10		Not Salient	4.15	0.07	4.05	0.89	1.11	.268	
				Salient	4.02	0.74	4.02	0.87	0.00	1.00	
	Hard work: Disadvantaged[b]	S10		Not Salient	4.19	0.06	4.20	0.76	0.03	.975	
				Salient	4.24	0.78	4.12	0.76	1.29	.198	
	Ability: Advantaged[b]	S10		Not Salient	4.15	0.07	4.05	0.87	1.04	.299	
				Salient	4.13	0.73	4.03	0.75	1.10	.273	
	Ability: Disadvantaged[b]	S10		Not Salient	4.21	0.06	4.14	0.82	0.70	.484	
				Salient	4.12	0.74	4.01	0.72	1.28	.202	
	Motivation: Advantaged[b]	S10		Not Salient	4.03	0.06	3.91	0.92	1.28	.202	
				Salient	4.01	0.73	3.90	0.85	1.17	.242	
	Motivation: Disadvantaged[b]	S10		Not Salient	4.09	0.07	4.04	0.83	0.50	.616	
				Salient	4.14	0.76	4.10	0.72	0.47	.636	
	Situational Attribution	2	.80	Standard	1.81	1.25	2.06	1.33	1.80	.072	†
		S5	.89	Policy-Narrow	2.86	1.59	2.96	1.52	0.73	.469	
				Policy-Abstract			3.17	1.61	2.29	.022	*
	Adversity	S3		Policy-Specific	3.08	1.09	2.68	1.26	3.45	<.001	***
				Policy-Range			2.41	1.14	6.15	<.001	***
	Academic Opport.: Advantaged[b]	S10		Not Salient	3.79	0.07	4.05	0.90	2.53	.012	*
				Salient	3.93	0.84	3.93	0.85	0.00	1.00	
	Academic Opport.: Disadvantaged[b]	S10		Not Salient	3.41	0.07	3.34	1.04	0.66	.513	
				Salient	3.19	0.96	3.22	0.93	0.25	.805	

(continued; 5 of 5)

Category	Label	Study	*r* / *α*	Conditions	Individual		Policy		*t*	*p*	
					M	*SD*	*M*	*SD*			
Causes of Applicant's Achievement (continued)	Household income: Advantaged[b]	S10		Not Salient	2.89	0.10	3.39	1.35	3.32	.001	**
				Salient	3.20	1.25	3.10	1.30	0.64	.523	
	Household income: Disadvantaged[b]	S10		Not Salient	2.19	0.07	2.30	1.05	0.96	.340	
				Salient	2.19	1.10	2.13	1.04	0.49	.625	
	Luck	S3		Policy-Specific	4.30	1.00	4.51	1.20	1.87	.063	†
				Policy-Range			4.48	1.38	1.44	.151	
	Luck: Advantaged[b]	S10		Not Salient	2.24	0.09	2.46	1.26	1.60	.11	
				Salient	2.47	1.27	2.32	1.16	1.05	.293	
	Luck: Disadvantaged[b]	S10		Not Salient	2.06	0.08	2.34	1.12	2.27	.024	*
				Salient	2.08	1.04	2.18	1.10	0.82	.415	

Note. This table compares non-choice measures for the Policy and Individual conditions. Comparisons are based on independent sample t-tests, except where indicated. The Category column indicates the general focus of the dependent variables. The Label column references the label given to the measure in the study methods; refer to the appropriate study (see the Study column) for more detail on the measure. The "r / α" indicates the correlation or reliability between items composing the measure (where relevant). The Conditions column indicates the specific Policy condition that is compared to the Individual condition; "Standard" refers to the Policy condition tested most frequently across studies (e.g., Study 1, 2, 3, 4 [Control]).

[a] This is a within-subjects measure. The t-statistic and p-value correspond to one-sample t-tests against the midpoint of the item's scale (0). Negative values indicate greater focus on this factor for the policy decision and positive values indicate greater focus on this factor for the individual decision.

[b] These measures are specific to the applicant stimuli shown in the study. Advantaged refers to the higher-scoring, advantaged applicant, whereas Disadvantaged refers to the lower-scoring, disadvantaged applicant.

† p < .10
** p < .05*
*** p < .01*
**** p < .001*

SM 4. OTHER ANALYSES

Construal level of text responses.

Throughout some of our studies (Studies 2, S01, S03-S06, S10, S12, and S13) we asked participants to explain why they selected the applicant and policy that they did. We used an online construal scoring tool (construalscore.com) to examine the construal level reflected in participants' responses. The authors of the paper explain:

> "In our paper, we calculated concreteness scores for each word in a set of online search queries using a dictionary developed by Paetzold and Specia (2016) and adapted from the Medical Research Council (MRC) machine readable dictionary (Coltheart 1981; Wilson 1988). This updated MRC dictionary contains 85,863 words that have been scored in terms of construal level. MRC concreteness scores typically range from 100 to 700, with lower numbers signifying a more abstract construal and higher numbers signifying a more concrete construal. We calculated the concreteness score of an entire search string (e.g., an online search query) by taking the average concreteness score of all words in the string that appeared in the MRC dictionary.
>
> We have developed this construal scoring tool as a web application so that academic researchers and marketing practitioners can determine concreteness scores for their own set of text strings. In addition to providing MRC concreteness scores for each string, our tool also outputs a concreteness score based on another dictionary of 39,954 words developed by Brysbaert et al. (2014). Brysbaert scores range from 1 to 5, with higher numbers indicating greater concreteness. For more information about the MRC and Brysbaert dictionaries, please consult the original research articles, listed in the References section below."

Throughout some of our studies (Studies 2, S01, S03-S06, S10, S12, and S13) we asked participants to explain why they selected the applicant and policy that they did. We used an online construal scoring tool (construalscore.com) to examine the construal level reflected in participants' responses. The authors of the paper explain:

We used the online tool to calculate MRC and Brysbaert scores for each participants' response. We omitted the words "applicant" and "policy" from the analyses so that the scores would not be artificially affected by language that was necessarily different between the conditions. We conducted t-tests on these score, comparing responses offered for policy decisions versus individual decisions. The MRC score for policy responses was significantly

lower (M = 315.44, SD = 19.57) than for individual responses (M = 318.95, SD = 21.56; t[3723] = 5.19, p < .001). The Brysbaert score for policy responses was also significantly lower (M = 2.29, SD = 0.25) than for individual responses (M = 2.36, SD = 2.36; t[3720] = 7.99, p < .001). Consistent with our theorizing, these results suggest that participants' explanations for their policy decisions reflected a more abstract construal than their explanations for individual decisions. While this finding is promising, more research is necessary to test whether construal level is indeed implicated as we theorize.

Moderation by individual-level characteristics. Across our studies we collected a number of individual-level demographic factors, namely, gender, age, race, income, education, and political attitudes. We did not find clear evidence that the policy-people gap is meaningfully moderated by any of these factors, with the exception of education. In the case of education, highly educated decision makers were more likely to favor the higher-scoring, higher-income applicant when choosing between individuals, leading to a larger policy-people gap. In interpreting these analyses, we relied not only on the statistical significance of the factors because some of these analyses include more than eight-thousand observations. Hence, for example, while age (p = .003, n = 8,639) significantly moderated the gap and gender marginally did so (p = .089, n = 8,559), we do not consider that they *meaningfully* moderate it.

Moderation by political orientation and values measures. Across 16 studies we measured participants' political orientation (a rough proxy for values) and specific values as conceptualized by different streams of research (i.e., egalitarianism, Ho et al., 2015; equality, Feldman, 1988; and the social justice moral motive, Janoff-Bulman & Carnes, 2016) and did not find that they predicted policy choices better than individual choices (see Table S9).

TABLE S9

Results of OLS Regressions Predicting Choice based on Decision Type and Values or Political Orientation (All Studies)

			Decision Type				Values / Political Orientation				Decision Type x Values / Pol. Or.			
Study	**Values Measure**	***n***	***b***	***CI*$_{95\%}$**		***p***	***b***	***CI*$_{95\%}$**		***p***	***b***	***CI*$_{95\%}$**		***p***
1	Political Orientation	165	0.31	0.11	0.51	.003	0.05	-0.04	0.13	.310	0.02	-0.09	0.14	.700
2	Political Orientation	347	0.06	-0.05	0.18	.292	0.09	0.04	0.14	<.001	-0.02	-0.09	0.04	.458
3	Political Orientation	790	0.12	0.05	0.18	.001	0.09	0.06	0.12	<.001	-0.02	-0.05	0.02	.401
4	Political Orientation	691	0.30	0.22	0.38	<.001	0.06	0.03	0.09	<.001	0.05	0.00	0.09	.030
S1	Political Orientation	1,977	0.32	0.27	0.36	<.001	0.02	0.01	0.04	0.011	.01	-0.02	0.03	.651
	Equality Preference	2,012	0.30	0.26	0.34	<.001	-0.10	-0.12	-0.08	<.001	0.03	0.00	0.06	.022
S2	Political Orientation	992	0.21	0.15	0.27	<.001	0.06	0.03	0.08	<.001	-0.03	-0.06	0.01	.132
S3	Political Orientation	601	0.35	0.27	0.43	<.001	0.07	0.04	0.09	<.001	-0.03	-0.08	0.01	.142
S4	Political Orientation	403	0.25	0.15	0.34	<.001	0.05	0.02	0.09	.005	0.00	-0.05	0.05	.948
S5	Political Orientation	559	0.16	0.08	0.25	<.001	0.05	0.02	0.09	.004	0.03	-0.03	0.08	.347
	SDO: Egalitarianism	576	0.11	0.04	0.18	.004	-0.12	-0.17	-0.07	<.001	-0.01	-0.07	0.06	.821
	SDO: Dominance	576	0.12	0.05	0.19	.002	0.11	0.06	0.16	<.001	-0.01	-0.09	0.06	.693
S6	Political Orientation	373	0.26	0.15	0.37	<.001	0.06	0.02	0.10	.004	-0.00	-0.06	0.05	.896
S7	Political Orientation	274	0.06	-0.09	0.21	.411	0.09	0.03	0.14	.002	-0.04	-0.11	0.03	.233
S8	Political Orientation	235	0.16	0.02	0.29	.027	0.05	-0.00	0.10	.054	0.02	-0.05	0.09	.549
S10	Political Orientation	301	0.23	0.12	0.34	<.001	0.08	0.03	0.12	<.001	-0.05	-0.11	0.01	.080
S11	Political Orientation	189	0.28	0.14	0.42	<.001	-0.02	-0.07	0.03	.427	0.05	-0.02	0.13	.154
S12	Political Orientation	298	0.06	-0.06	0.18	.301	0.04	-0.01	0.08	.111	-0.03	-0.09	0.03	.331
S13	Political Orientation	285	0.21	0.09	0.33	.001	0.04	-0.01	0.09	.096	-0.03	-0.09	0.04	.386
	Social Justice Moral Motive	287	0.22	0.11	0.33	<.001	-0.13	-0.19	-0.08	<.001	0.05	-0.04	0.13	.279

Note. We conducted OLS regressions on the choice (1 = Higher-scoring, advantaged option, 0 = Lower-scoring, disadvantaged option) with factors for Decision Type (1 = Individual, 0 = Policy), Values (1-7, indicating greater agreement with scale) or Political Orientation (1 = Ext. Liberal, 4 = Moderate, 7 = Ext. Conservative), and Decision Type x Values / Political Orientation. For ease of interpretation, we centered the political orientation variable at its midpoint and the values variables at their means. We examine only the standard Individual versus Policy conditions (e.g., Study S4's Policy-Abstract condition) within each study.

SM 5. SUPPLEMENTAL REFERENCES

Brockner, J., Chen, Y. R., Mannix, E. A., Leung, K., & Skarlicki, D. P. 2000. Culture and procedural fairness: When the effects of what you do depend on how you do it. ***Administrative Science Quarterly***, 45*:* 138-159.

Castilla, E.J. & Benard, S. 2010. The paradox of meritocracy in organizations. ***Administrative Science Quarterly***, 55: 543-676.

Critcher, C. R., & Dunning, D. 2013. Predicting persons' versus a person's goodness: Behavioral forecasts diverge for individuals versus populations. ***Journal of Personality and Social Psychology***, 104: 28–44.

Feldman, S. 1988. Structure and consistency in public opinion: The role of core beliefs and values. ***American Journal of Political Science***, 32: 416–440.

Graham, J., Haidt, J. & Nosek, B.A. 2009. Liberals and conservatives rely on different sets of moral foundations. ***Journal of Personality and Social Psychology***, 96: 1029-1046.

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 SDO_7 scale. ***Journal of Personality and Social Psychology***, 109: 1003-1028.

Hughes, J. S. 2017. In a moral dilemma, choose the one you love: Impartial actors are seen as less moral than partial ones. ***British Journal of Social Psychology***, 56: 561-577.

Humphreys, A., Isaac, M. S., & Wang, R. J. H. 2021. Construal matching in online search: applying text analysis to illuminate the consumer decision journey. ***Journal of Marketing Research***, 58: 1101-1119.

Janoff-Bulman, R., & Carnes, N. C. 2016. Social justice and social order: Binding moralities across the political spectrum. ***PLoS One***, 11.

Kreps, T. A., & Monin, B. 2014. Core values versus common sense: Consequentialist views appear less rooted in morality. ***Personality and Social Psychology Bulletin***, 40: 1529-1542.

Lucas, B.J. & Kteily, N.S. 2018. (Anti-) egalitarianism differentially predicts empathy for members of advantaged versus disadvantaged groups. ***Journal of Personality and Social Psychology***, 114: 665-692.

Leventhal, G. S. 1976. The distribution of rewards and resources in groups and organizations. In L. Berkowitz & W. Walster (Eds.), ***Advances in experimental social psychology***, Vol. 9,

91-131. New York, NY: Academic Press.

Leventhal, G. S. 1980. What should be done with equity theory? New approaches to the study of fairness in social relationships. In K. Gergen, M. Greenberg, & R. Willids (Eds.), ***Social exchange: Advances in theory and research***, 27-55. New York, NY: Plenum Press.

Lind, E. A., Tyler, T. R., & Huo, Y. J. 1997. Procedural context and culture: Variation in the antecedents of procedural justice judgments. ***Journal of Personality and Social Psychology***, 73: 767-780.

Luguri, J.B., Napier, J.L., & Dovidio, J.F. 2012. Reconstruing intolerance: Abstract thinking reduces conservatives' prejudice against nonnormative groups. ***Psychological Science***, 23: 756-763.

Mahfud, Y., Badea, C., Verkuyten, M. & Reynolds, K. 2018. Multiculturalism and attitudes toward immigrants: The impact of perceived cultural distance. ***Journal of Cross-Cultural Psychology***, 49: 945-958.

McClanahan, K. J., Ho, A. K., & Kteily, N. S. 2019. Which group to credit (and blame)? Whites make attributions about White-minority biracials' successes and failures based on their own (anti-) egalitarianism and ethnic identification. ***Group Processes & Intergroup Relations***, 22: 631-654.

Napier, J.L., & Luguri, J.B. 2013. Moral mind-sets: Abstract thinking increases a preference for "individualizing" over "binding" moral foundations. ***Social Psychological and Personality Science***, 4: 754-759.

O'Laughlin, M. J., & Malle, B. F. 2002. How people explain actions performed by groups and individuals. ***Journal of Personality and Social Psychology***, 82: 33–48.

Reynolds, S.J., Leavitt, K. & DeCelles, K.A. 2010. Automatic ethics: The effects of implicit assumptions and contextual cues on moral behavior. ***Journal of Applied Psychology***, 95: 752-760.

Russell, D. W. 2002. In search of underlying dimensions: The use (and abuse) of factor analysis in Personality and Social Psychology Bulletin. ***Personality and Social Psychology Bulletin***, 28: 1629-1646.

Shrout, P. E., & Bolger, N. 2002. Mediation in experimental and nonexperimental studies: new procedures and recommendations. ***Psychological Methods***, 7: 422-445.

Sinclair, R. C., & Mark, M. M. 1991. Mood and the endorsement of egalitarian macrojustice

versus equity-based microjustice principles. ***Personality and Social Psychology Bulletin***, 17: 369-375.

Skitka, L. J., & Morgan, G. S. 2014. The social and political implications of moral conviction. ***Political Psychology***, 35: 95-110.

Thibaut, J. W., & Walker, L. 1975. ***Procedural justice: A psychological analysis***. Hillsdale, NJ: Erlbaum.

Uhlmann, E. L., Zhu, L. L., & Tannenbaum, D. 2013. When it takes a bad person to do the right thing. **Cognition**, 126: 326-334.

Wright, J.C. & Baril, G. 2011. The role of cognitive resources in determining our moral intuitions: Are we all liberals at heart?. ***Journal of Experimental Social Psychology***, 47: 1007-1012.

Yogeeswaran, K. & Dasgupta, N. 2014. The devil is in the details: Abstract versus concrete construals of multiculturalism differentially impact intergroup relations. ***Journal of Personality and Social Psychology***, 106: 772-789.

Zdaniuk, A., & Bobocel, D. R. 2011. Independent self-construal and opposition to affirmative action: The role of microjustice and macrojustice preferences. ***Social Justice Research***, 24: 341-364.