Motivated moral judgments about freedom of speech are constrained by a need to appear consistent

Nikolai Haahjem Eftedal^{a*}

Lotte Thomsen^{a,b}

^aDepartment of Psychology, University of Oslo, Oslo, Norway

^bDepartment of Political Science, Aarhus University, Aarhus, Denmark

The research was funded by grants 0602-01839B and 231157/F10 from the Danish- and Norwegian Research Councils, respectively (to L.T.).

*Corresponding author

E-mail: n.h.eftedal@psykologi.uio.no

Author contributions: N.H.E & L.T. designed research; N.H.E collected and analyzed data; N.H.E. & L.T. wrote and edited the paper

Abstract:

Speech is a critical means of negotiating political, adaptive interests in human society. Prior research on motivated political cognition has found that support for freedom of speech depends on whether one agrees with its ideological content. However, it remains unclear if people (A) openly hold that some speech should be more free than other speech; or (B) want to appear as if speech content does not affect their judgments. Here, we find support for (B) over (A), using social dominance orientation and political alignment to predict support for speech. Study 1 demonstrates that if people have previously judged restrictions of speech which they oppose, they are less harsh in condemning restrictions of speech which they support, and vice versa. Study 2 finds that when participants judge two versions of the same scenario, with only the ideological direction of speech being reversed, their answers are strongly affected by the ordering of conditions: While the first judgment is made in accordance with one's political attitudes, the second opposing judgment is made so as to remain consistent with the first. Study 3 is a preregistered replication and elaboration on Study 2. Study 4 suggests that this effort to appear consistent functions to oblige the opposition to (also) protect one's own speech: We find that support for equal protection of all speech is stronger if a member of the opposition proposes it, rather than a member of one's own political coalition. These results are consistent with notions of an evolutionary arms race of social manipulation, and suggest that although people selectively endorse universal moral principles depending on their own political agenda, they conceal their bias from others and, perhaps, themselves.

Keywords: motivated reasoning, moral judgment, freedom of speech, self-deception, social dominance, political ideology

Motivated moral judgments about freedom of speech are constrained by a need to appear consistent

1. Introduction

“Goebbels was in favor of free speech for views he liked. So was Stalin. If you’re really in favor of free speech, then you’re in favor of freedom of speech for precisely the views you despise. Otherwise, you’re not in favor of free speech.” – Noam Chomsky

“If we continue to teach about tolerance and intolerance instead of good and evil, we will end up with tolerance of evil” – Dennis Prager

Someone spreading controversial beliefs or attitudes is likely to face consequences. However, most societies have rules, both written and informal, which regulate how harsh these consequences can be. These rules will likely not prevent people from countering a message through civil and honest presentations of arguments and evidence. But they could, to some extent, restrain people from hurting others socially, financially, or physically, or from taking away their platforms. Even people who despise a view that is being spread might object if they feel others are breaking these rules for how to legitimately counter it. (We will refer to these rules for how to counter speech as *speech norms*.)

The world is currently undergoing a revolution in information technology, with the internet and social media expanding most people’s audiences by orders of magnitude. Norms for how information should be spread and restrained are thus increasingly relevant. As highlighted by recent controversies (see e.g. TheFire, 2019; Bey, 2017), there is no clear consensus on what exactly these norms should say. For example, some hold that the removal of certain platforms is legitimate (e.g., Munroe, 2014), while others disagree (e.g. Campbell, 2018). Similarly, attempts at getting people fired for their views are applauded by some (e.g. Bey) and decried by others (e.g., Crook, 2017).

A central question is whether speech content should be morally relevant when making judgments about tactics to counter speech. For instance, does the blameworthiness of sabotaging a speaking event depend on the message being delivered there? The recurring finding from experimental psychological studies examining these kinds of judgments (Lindner & Nosek, 2009; Crawford, 2014; Brandt, et al., 2015; Crawford & Pilanski, 2015; Crawford & Xhambazi, 2015; White II & Crandall, 2017) is that participants indeed seem to take speech content into account: They judge an action to be more permissible when used against speech they oppose than when used against speech they support. To illustrate, Crawford (2014) found that conservatives were more inclined to agree that abortion activists should be banned from distributing pamphlets and fliers on a campus when described as pro-choice rather than pro-life, while liberals showed the opposite pattern.

However, it remains unclear from prior studies if

(A) Participants are explicitly endorsing *content-sensitive speech norms* saying that the blameworthiness of actions to counter speech depends on what that speech is. If so, prior results simply reflect participants' honest opinions that certain views should have less protection than others.

(B) Participants seek to appear as if they support *content-neutral speech norms* saying that actions to counter speech are equally blameworthy, regardless of what is being said. If so, prior results reflect selective, context-specific adjustments of this general level of blame for restricting speech: If a participant supports/opposes the speech in question, they will take the stance that it is *always* wrong/right to counter speech in the specific way they just evaluated.

Here, we seek to disambiguate between these two competing interpretations. To the extent that prior studies have interpreted results to be due to either (A) or (B), (A) - that people are openly content-sensitive - has mostly been emphasized. And content-sensitive speech norms have certainly existed in many civilizations throughout history (e.g., against blasphemy), both as instantiated through laws and through social sanctions (Green & Karolides, 2014; Graham, 2004). Yet, we favor interpretation B. As argued by e.g. Descioli and Kurzban (2013), a common feature of moral systems is that rules should be universal and equally applied across cases, with only a minimum of caveats and exceptions. This is a sentiment that has deep evolutionary- and developmental roots (Boehm, 1999, see Thomsen, 2020 for review). At the same time, there is a constant tension in morality between cooperation for the common good and defection to serve other interests (Wilson, 1998), which might have led to an evolutionary arms race in the domain of social negotiation and manipulation (Dawkins & Krebs, 1979; Trivers, 1971). Applied to the context of speech protection, the winning strategy could then be to *appear as if* applying rules equally across all kinds of speech while actually making judgments in a partisan manner, to the extent that this can be plausibly denied. Thus, while most people prefer some views to win out over others (as shown in prior research), we predict that they will still want to see and portray themselves as supporting norms that give equal protections to all (or at least most) views. Before further elaborating on this prediction and how we test it, we will first discuss the value of being able to shape what information other people are-, and are not-, exposed to.

1.1 Ideology

Views affect how people behave, and which policies and norms they put in place. Whenever people act in ways that make things worse, or fail to act in ways that make things better, they might have acted differently if their views were different. So, from a purely impartial standpoint, there are very good reasons to care what other people think: Changing people's minds can improve the world. And preventing minds from being changed can keep the world from deteriorating. (Note that sorting views into good and bad is different from sorting into true and false; see e.g. Boström, 2011)

Interestingly, people may have quite contrasting opinions on which views should and should not be spread. This could reflect honest disagreements about how to best reach goals we all share. It could also reflect conflicts of interests when our goals are not fully shared (see Alexander, 2018): Altering other people's views can serve selfish and parochial goals as well as altruistic and impartial ones. And the beliefs and values that would make the world better on the whole are not necessarily the same as those that serve your personal preferences: The world is full of trade-offs where improving conditions for some often means making things worse for others.

Societies typically have several ideological groups working to shift trade-offs to their benefit. As implied by the word "ideological", this work involves influencing the flow of ideas (Haidt, 2012; Greene, 2014). Being accepted into such groups has benefits, and so taking part in these actions to spread or restrict views may also be motivated by a desire to signal one's loyalty to the group in question (Kahan, 2014, cf. Greene, 2014).

However, which groups one seeks acceptance from in the first place likely reflects the degree to which their ideological goals overlap with one's -own (cf. e.g., Blumer, 1958; LeVine & Campbell, 1972; Sidanius & Pratto, 1999; Tooby & Cosmides, 2010).

There are typically many more conflicts of interests in a society than there are ideological groups. Groups will tend to have multiple agendas at once, and each single member might not care equally about all of these. Social Dominance Theory addresses one of the fundamental, general ideological conflicts, around which groups tend to form and legitimize their agenda, namely that of societal hierarchy versus equality.

1.1.1 Social Dominance Orientation and legitimizing myths

Social Dominance Theory suggests that the root of many conflicts in moral and ideological matters are the dominance hierarchies that are ubiquitous across human, surplus-producing societies and individual differences in the motives to either sustain or attenuate them (Sidanius, 1993; Sidanius & Pratto, 2004; Sidanius et al, 2016; Kleppesø et al, 2019; cf. Marx, 1972). Indeed, philosophers from Thucydides and Hobbes to Foucault posit that politics is undergirded by the question of who will dominate whom (McClelland, 1996), and social dominance is represented and motivates affiliation even among infants (Thomsen, et al., 2011; Thomas et al., 2018). The Social Dominance Orientation (SDO) scale (Pratto et al., 1994; Ho, et al., 2012; Ho, et al., 2015) captures individual differences in preferences for intergroup hierarchies.

Social Dominance Theory further suggests that motives regarding hierarchies - whether one wants to sustain or attenuate them - lead to the production and perpetuation of *legitimizing myths*. Legitimizing myths are narratives and truth-claims that serve to justify one's social and political preferences. High SDO is associated with support for hierarchy-enhancing legitimizing myths (Ho et al., 2012; 2015; Lucas & Kteily, 2018; Thomsen et al., 2010) which function to justify hierarchies, be they true or false. Examples include the belief that the current distribution of status and power in

society has come about through meritocratic processes, or that negative stereotypes about low status groups are true. Conversely, low SDO predicts support for hierarchy-attenuating legitimizing myths, such as feminism and anti-racism, which discredit the validity of intergroup hierarchies.

The idea that people seek to “legitimize” their preferences is echoed across the literature on motivated reasoning (Kunda, 1990; Von Hippel & Trivers, 2011; Mercier & Sperber, 2011; Simler & Hanson, 2017; Ditto et al., 2018). While one’s true motives might be to serve a certain agenda, people generally seek to present themselves as caring mostly about the truth and the common good. Being open about having biased reasons for trying to influence others will typically make one less persuasive, while appearing to be impartial and honest can result in a better shot at getting them on your side. In this sense, one might argue that politics is largely about persuading others that general rules that happen to benefit one’s specific interests also serve the greater common good, thus appealing to everyone for support (Petersen, 2015). Trivers further argues that it is in many ways useful to not only deceive others about one’s strategic motives, but to keep one’s own conscious mind in the dark as well; to be “self-deceived” (Trivers, 2000; 2011). Being self-deceived is less cognitively demanding than being dishonest; it reduces risks of being judged as a liar; and it helps you present a more convincing case for your position.

1.2 Appearing impartial

A preference for selective protection of certain kinds of speech is perhaps the kind of preference that remains mostly hidden from the conscious mind. It is difficult to legitimize such a policy to an audience that includes people who support viewpoints that would then be protected *less*.

People could then instead portray themselves as supporting largely content-neutral speech norms, where a wide range of views are equally protected. Such norms are what you would prefer your ideological opposition to support. Content-sensitive norms saying that “bad views should have less protection” are problematic if practiced by people who disagree with you about which views are good and bad. They might use these norms to justify silencing views you support and protecting views you oppose. Endorsing shared principles of equal protection is a way to compromise.

Notably, holding that the legitimacy of a tactic is independent of the speech it is used against will still permit judging people based on when they prefer to use this tactic. You may for example at the same time see social shaming as a legitimate tactic and also condemn someone for shaming the expression of a view you support. But this condemnation is then directed at the person’s dislike of a view you like, rather than at their use of shaming *per se*.

There are reasons to prefer content-neutral speech norms also among your ideological allies, not just your opposition. John Stuart Mill voiced several such reasons in “On Liberty” (1859/1966), which have since been echoed by others (Galef, 2018; Alexander, 2014a). For example, Mill (p. 24) argues that silencing wrongful views deprives people of “the clearer perception and livelier impression of truth, produced by its collision with error.” Additionally, content-neutral norms could

function as a safe-guard against the eventuality that power might at some point fall into the wrong hands (Brennan, 2015), or that suppressed views might “go underground” only to emerge again later in a more destructive form (Pinker, 2007). Nevertheless, it is apparent from the existing experimental record that many people have not been swayed into adopting content-neutral norms from arguments such as these: Otherwise they would not have systematically judged restrictions of speech they oppose to be more permissible than restrictions of speech they support.

Importantly, it cannot be conclusively shown from prior between-subject studies that any one specific participant was in fact being content-sensitive in their judgments; this can only be inferred when looking at all responses in conjunction. For each response, participants can claim that they would indeed have responded exactly the same had the ideological content of speech been reversed. This type of psychological wiggle room is common in the real world as well; no two situations are exactly alike, so you can usually point to unique features unrelated to speech content to justify judging in ways that benefit your own position. As long as you can maintain plausible deniability, you can get away with judging selectively while still publicly endorsing content-neutral norms. Such a strategy is only effective, however, so long as its double standard is not laid bare - for others, nor perhaps even yourself (Trivers, 2011).

1.3 Present Study

Here, we examine whether people seek to hide their tendency to give more freedom to speech they support (a tendency that has been demonstrated in prior research and which we replicate here). Specifically, we investigate effects of altering the ordering of the speech participants make judgments about. In Study 1, we measure participants’ SDO and ask them to respond to scenarios depicting restrictions of speech, where speech content is manipulated to serve goals related to having either high or low SDO. We then analyze whether participants become more supportive of rights for speech they dislike if they have previously judged speech they support, and vice versa. A façade of neutrality is harder to maintain if certain types of speech are systematically given less protection than other types of speech, and so we predict that the bias in free speech judgments shown in the prior literature should be dampened by having first made judgments for the opposite kind of speech.

In Study 2, we manipulate the ideology of speech *within subjects*, such that participants first judge a scenario with one type of speech, and then judge the *exact same* scenario except that the ideology of speech content is reversed. Participants are unaware that the second scenario is coming when responding to the first, and they are unable to alter their responses to the first upon seeing the second. Under hypothesis A above, where participants openly endorse that some speech should be less protected than other speech (“content-sensitive speech norms”), the order of presentation should not matter: Participants should simply judge speech suppression differently depending on speech content. On the other hand, if B is true and participants seek to maintain the narrative that they are content-neutral in their judgments, then the ordering of scenarios will affect responses to both of

them. By having scenarios be exactly matched on everything but the ideology of speech content, participants lose the ability to plausibly claim that content-sensitive judgments are in fact content-neutral. To maintain the appearance of content-neutrality, responses to the second scenario must be similar to those for the first. This means that the ordering of conditions will affect both responses, so that either both become more lenient (if the speech in the first scenario was opposed by the participants), or both become more harsh (if the speech in the first scenario was supported). See Figure 1 for an illustration of these kinds of order effects.

Study 3 is a pre-registered replication of Study 2, which also explores the potential moderating roles of introspectiveness and two measures of moral integrity. Both introspectiveness and integrity could dampen the tendency to shift judgments depending on speech content, if it is indeed the case that this bias is the result of processes that are on the fringes of conscious access (and thus more accessible to the introspective) and which are morally questionable (such that people with more integrity will be less likely to follow them).

Finally, in Study 4 we investigate if the expressions of support for content-neutral speech norms are themselves products of politically motivated reasoning. Specifically, we investigate if support for content neutral norms is heightened if they are proposed by one's ideological opposition. This hypothesis follows from the idea that the reason people seek to appear content-neutral is their hope that their opposition will return the favor. Again, we manipulate which ideological side proposes content-neutral norms within subjects, predicting an order effect reflecting that people seek to seem as if their endorsements of speech norms are independent from ideological considerations.

Figure 1. Illustration of order effects.

2. Study 1

In Study 1 we use the SDO-7 scale (Ho et al., 2015) to predict participants' judgments of three scenarios describing the use of different tactics for restraining the spread of views: online shaming, selective standards of evidence (Alexander, 2014a), and job-firing. We randomly vary the ideological speech content affected by these tactics *between subjects*, to either support goals related to SDO (Pro-SDO condition) or to oppose them (Contra-SDO condition). In order to investigate our two competing hypotheses for interpretation of prior studies, we analyze the effects of prior exposure to scenarios in the opposite ideological condition to the current one, using a randomization scheme which makes this approach viable. Insofar as participants seek to seem consistent in judgments regardless of speech content, exposure to the opposite condition should dampen the selective protection of favored speech observed in prior studies.

2.1 Participants

We recruited a sample of 300 participants from Amazon's Mechanical Turk, ensuring ~90% power to detect effects of similar magnitude to those observed in prior studies. All power analyses are done using G'Power 3.1 (Faul, Erdfelder, Lang, & Buchner, 2007).

We excluded 53 participants who failed more than one of our three attention checks. Among the remaining 247 participants ($M_{age} = 37.1$, $SD_{age} = 11.3$, 104 females), there were 69% Caucasians, 18% Asian American, 9% African Americans, 2% Hispanics, 1% Native Americans, and 1% Jews. This project complies with, and was conducted in accordance with, ethical standards as outlined by the Norwegian Centre for Research Data.

2.2 Procedure

Participants gave an informed consent and basic demographic information before filling out the short 8-item version of the SDO-7 scale ($\alpha = .93$, $M = 2.39$ out of 7, $SD = 1.31$). They then responded to three questions for each vignette about the legitimacy of the tactic used to counter speech. Vignettes were presented in random order. The speech in the vignettes supported goals associated with either high- or low SDO, depending on the experimental condition. Except for minor adjustments, the three vignettes match the corresponding scenarios in Studies 2 and 3, which are displayed in Box 1 (labeled "Job security", "Academic freedom", and "Online hate"; for the verbatim texts used in Study 1, see the Online Supplement).

For each vignette, participants' responses to the associated statements were averaged (with reversed scoring of reversed items) to form speech protection measures serving as the dependent variables for our analyses ($\alpha's > .68$).

For 166 of the 247 participants (200 of the original 300), we randomized the experimental condition separately for each vignette, so that they were likely to receive some vignettes from each condition. For the remaining 81 participants (100 out of original 300), we rather randomized condition for all vignettes simultaneously, so that all vignettes were in the same condition. Following this randomization scheme, approximately half of responses to each individual vignette would be preceded by responses to another vignette in the other experimental condition (i.e., framed in the opposite ideological direction).

2.3 Results

All analyses and figures across all four studies were done in the statistical computing environment R (R Core Team, 2017). The regression models to investigate selective support of free speech can be seen in Table 1. The significant interactions between SDO and Condition in all three vignettes are consistent with prior work, suggesting that participants indeed support free speech norms selectively, depending on their attitudes to the speech content: They condemn restriction of speech more for content which they support and less for content with which they disagree (see Table 1).

The three-way interaction between SDO, Speech Content and Exposure to Both Conditions (EBC) indicates whether this selective protection-effect changes when the respondent faces an impartiality issue because s/he has prior exposure to vignettes framed in the opposite ideological direction. Combining all three vignettes in a mixed effects model, with random coefficients to account for dependencies between multiple observations from the same subject and from the same vignette, the three-way interaction between SDO, Speech Content and EBC is significant (see Table 2). The negative coefficient is about half in magnitude but with opposite sign to the interaction between SDO and Speech Content. This demonstrates that the motivated, selective protection of speech with which one agrees is quite substantially diminished by prior exposure to speech framed in the opposite ideological direction about a different case: If people have previously evaluated a case of restricting speech which they support, they will more strongly condemn a case of restricting speech which they oppose. Conversely, if people have previously evaluated a case of restricting speech which they oppose, they will condemn other cases of restricting speech which they support less strongly.

Box 1. Vignettes in Study 2 and 3

Job security

Stanley works as an accountant for a small firm. In his free time, he is active on various political messaging boards online. During the election, to express his disdain for [Hillary Clinton/Donald Trump], he posted demeaning and obscene images with [Clinton's/Trump's] face crudely photoshopped into them. Though he posted this under a pseudonym, offended [Clinton/Trump] supporters were able to ascertain his identity. They alerted his employer, who was also a [Clinton/Trump] supporter, to the posts. The employer then decided to fire Stanley.

1. Stanley's employer should have the right to fire him, even if the reason is unrelated to the quality of work produced. [Reversed]
2. Stanley's right to freedom of expression was compromised here.
3. The people contacting Stanley's employer were morally wrong in doing that.
4. There are limits to freedom of expression, and Stanley might have stepped over those limits here. [Reversed]

Online hate

Tina made a long post on Facebook, thinking only her small circle of friends would see it. She argued there is currently too much [sensitivity/insensitivity] around certain topics: People are too [often/rarely] called out for things they've said about women or minorities, and faced with social disapproval. Tina urged the readers of her post to become less [eager/reluctant] to apply labels such as "racist" and "sexist" in their daily lives.

The post was well liked among Tina's friends, many of whom circulated the post further. Eventually, however, it was seen by a prominent blogger, who decided to criticize and ridicule it on his blog. Following this, Tina received a torrent of critical and abusive messages from the blog's readers over a period of about a week, causing her great distress.

1. Ideally, Tina should have been free to speak her mind without fear of a backlash.
2. Tina's critics should be more tolerant of differences in opinion.
3. Facebook is a public forum, so Tina should have been prepared to face the consequences of her post. [Reversed]
4. The prominent blogger was in the wrong for bringing negative attention to Tina

Note. All four vignettes used in Study 2 and 3, with "headlines" (not seen by participants) indicating the topic raised. Bold font and brackets indicate points were the two experimental conditions differ, with the text for the Pro-SDO condition first, and then the text for the Contra-SDO condition after the slash. The numbered points are the statements participants rated their agreements with, on 7-point likert scales.

No-platforming

At a university, a student union had decided to invite a controversial speaker to an event they were organizing. The speaker was a prominent [critic of/advocate for] multiculturalism and open borders, who has argued that western culture is [superior/not superior] to others, and that an influx of non-westerners into the US would be [a detriment/an enrichment].

As a response to the invitation, outraged students held loud and disruptive protests and they threatened to sabotage the event through measures such as interrupting the speaker with megaphones. The speaker was then disinvited by the university administration, in compliance with the protesters' demands.

1. The speaker's views are sufficiently harmful that it is legitimate to silence them. [Reversed]
2. The protesters are only exercising their free speech; they're not doing anything wrong here. [Reversed]
3. The university administration should not be blamed for disinviting the speaker, as it is fully within its rights to do so. [Reversed]
4. It is wrong for the protesters to try to remove the speaker, regardless of how abhorrent they find his views.

Academic freedom

Beth is a psychology researcher studying personality and individual differences. She is particularly interested in studying the ways democrats and republicans differ psychologically. She is a strongly committed [conservative/progressive] herself. Among the conclusions from her studies are that [democrats/republicans] tend to score higher on psychoticism and authoritarianism, and that they have a higher tendency to falsely report having socially desirable characteristics.

Some [progressive/conservative] researchers, who are uncomfortable with Beth's conclusions, have put her work under very harsh scrutiny. They have also criticized Beth for pursuing research questions that only serve to increase division, and they have implied that she is biased and lacks scientific integrity.

1. Beth's conclusions are controversial, so increased scrutiny is justified. [Reversed]
2. Beth's chosen topic of research is best left unexplored. [Reversed]
3. The critics are stifling Beth's academic freedom.
4. Beth should keep her findings out of lectures and writings, as they are so divisive. [Reversed]

Table 1. The coefficients from the regression models of SDO, Condition, and their interaction on speech protection, across three scenarios.

	<u>Online Shaming</u>			<u>Academic Freedom</u>			<u>Job Security</u>		
	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		3.70	[3.53, 3.86]		5.39	[5.22, 5.56]		4.96	[4.70, 5.21]
SDO	.11	.10	[-.07, .28]	-.22*	-.16	[-.29, -.02]	-.08	-.09	[-.28, .10]
SpeechContent	-.02	-.02	[-.19, .14]	.06	.05	[-.12, .23]	-.04	-.05	[-.30, .20]
EBC	.01	.01	[-.23, .25]	.04	.07	[-.16, .30]	-.04	-.10	[-.46, .25]
SDO*SpeechContent	.21*	.19	[.02, .37]	.19*	.13	[.00, .27]	.27**	.29	[.10, .48]
SDO*EBC	-.14	-.18	[-.42, .06]	-.05	-.05	[-.22, .13]	-.08	-.12	[-.38, .15]
SpeechContent*EBC	.05	.07	[-.17, .31]	-.09	-.12	[-.35, .12]	-.03	-.05	[-.41, .30]
SDO*SpeechContent*EBC	-.09	-.12	[-.36, .12]	-.10	-.09	[-.27, .08]	-.11	-.16	[-.43, .10]

Note. SDO is a participant's mean centered score on the Social Dominance Orientation scale. SpeechContent is contrast coded as 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. EBC indicates whether a participant has been exposed to both experimental conditions at the time of responding, coded as 1 when this is the case and 0 when it is not. Due to how we randomized conditions (described in Methods), this was the case for approximately half of the responses to each vignette. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

Table 2. Fixed effects from a mixed effects regression model of SDO, Condition, EBC, and their interactions on judgments, on pooled data for the three vignettes. The model has random effects accounting for dependencies between multiple responses across subjects and vignettes.

	β	<i>Pooled data</i>	
		<i>b</i>	<i>CI</i> _{95%}
(Intercept)		4.66	[3.66, 5.66]
SDO	-.08	-.08	[-.16, .01]
SpeechContent	-.00	-.00	[-.12, .11]
EBC	.01	.03	[-.13, .20]
SDO*SpeechContent	.21***	.21	[.12, .29]
SDO*EBC	-.05	-.06	[-.19, .06]
SpeechContent*EBC	-.02	-.03	[-.19, .13]
SDO*SpeechContent*EBC	-.11**	-.15	[-.27, -.03]

Note. SDO is a participant's mean centered score on the Social Dominance Orientation scale. SpeechContent is contrast coded as 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. EBC indicates whether a participant has been exposed to both experimental conditions at the time of responding, coded as 1 when this is the case and 0 when it is not. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$. Degrees of freedom for significance tests in all our mixed models are calculated using the Kenward-Roger approximation (Kenward & Roger, 1997).

3. Study 2

In Study 1, participants can still plausibly deny any influence of speech content on judgments, even if they are exposed to both ideological flavors of speech across the different vignettes: The scenarios have other differences besides the speech content, which could justify differences in judgments. To directly block such avenues for plausible deniability, in Study 2 we manipulate speech *within* subjects, so that all participants respond to both versions of each vignette.

We randomized the order of presentation *between* subjects: Participants either first make judgments about four vignettes with speech they support and then the same vignettes with speech they oppose, or the other way around. Importantly, participants are unaware that the second set of vignettes is coming when answering the first set, and they are also unable to go back to change the answers to the first vignettes after having seen the second set.

If participants openly support content-sensitive speech norms, then the order of presentation should not matter: Participants should simply judge differently depending on speech content in each condition, since they are open about thinking that the extent to which an action is a blameworthy violation of speech rights depends on the content of the speech that is counteracted.

If participants are instead motivated to portray themselves as being content-neutral in their judgments, as we predict, then they should be motivated to *answer similarly* to their replies to the first vignettes when answering their ideological opposites the second time. For instance, if a participant first judges that it is blameworthy to get someone fired for having a view she supports, she might feel compelled to then say that it is equally blameworthy to get someone fired for the exact opposite view, in order to maintain consistency. But if ordering was switched, so that the first judgment was about the firing of someone with views she opposed, then both judgments would have been more lenient.

This prediction is contingent not only upon participants' motivation to push for content-neutral speech norms, but also on them truly believing that they support such norms. Given that prior findings of selective protection of speech hold, such a belief would then require some level of self-deception. Without this caveat, there would be little reason to expect participants to seek to maintain the impression of consistency even when filling out an anonymous online survey where they face no penalty for moral inconsistency.

3.1 Participants

As we expected at least as large effects in Study 2 as those in Study 1, we once again recruited 300 participants from MTurk. Since the distribution of SDO scores among our sample in Study 1 had a clear overrepresentation of low scores, Study 2 recruited participants that were 50-50 Democrats and Republicans, by utilizing MTurks filtering function. As SDO is robustly correlated with political affiliation (Ho et al., 2012; 2015), this should lead to substantially more variability in SDO scores, which then increases our ability to detect interaction effects involving SDO.

We excluded participants on a per-vignette basis if they did not pass an attention check for each of the vignettes. We also excluded participants who did not clear at least 3 of the 4 attention checks from all analyses. This yielded a mean sample size of 244 participants per analysis.

Among the 262 participants not fully excluded from the study ($M_{\text{age}} = 40.0$, $SD_{\text{age}} = 12.1$, 141 females), there were 84% Caucasians, 5% Asians, 6% African Americans, and 4% Hispanics.

3.2 Procedure

Participants gave an informed consent and basic demographic information before filling out the short 8-item version of the SDO-7 scale ($\alpha = .93$, $M = 2.64 \pm SD$ of 1.50). They then made moral judgments about four different free speech vignettes (in random order) where the speech being suppressed was either pro- or contra goals associated with high SDO.

The four vignettes used are shown in Box 1. They resemble those in Study 1, with the addition of a fourth scenario about a public speaking event being sabotaged (“No-platforming”). Participants’s average responses across the four judgments for each vignette were the dependent variables in analyses. Responses were always scored such that higher scores indicate more support for the speaker. Cronbach’s α was at $>.76$ for three of the four vignettes. The exception was the *Academic Freedom* vignette, with $\alpha=.57$. We report analyses where these items are nonetheless combined into a single scale, since the pattern of results is robust across all four single items.

After responding to the first set of vignettes, participants were given attention checks and then asked to respond to the scenarios once again, but now with alterations. They indicated the extent to which they felt the changes to the text were relevant to the associated judgments (on a 5-point likert from “not at all” (1) to “to a large extent” (5)), before they rated agreement to the judgments once again with the new scenario in mind. They had no option to go back and view or alter previous responses at this stage. The ordering of vignettes was the same as for the first set.

3.3 Results

When considering only responses to the first set of vignettes, there are a significantly positive coefficients for the interactions between SDO and Condition ($p's < .001$), for all four vignettes. This replicates the prior finding that people favor increased protection for views they support (See Table 3 and the first column of Figure 2 for results and visualizations). Low SDO predicts increased protection of the speech in the Contra-SDO condition and vice versa for high SDO.

For the *Job Security* vignette only, there is a negative main effect of Condition ($p < .01$), meaning that the posting of photoshopped images of Hillary Clinton was tolerated less than photoshopped images of Trump, across all participants.

For the *No-platforming* vignette only, there is a positive main effect of SDO ($p < .001$), indicating that the average protection of speech across both conditions was higher among those with high SDO here.

The second set of responses is towards the same vignettes and statements as in the first set, except that all the judgments have been preceded by making the same judgments in the opposite ideological condition. As predicted, when running the same analysis as above on these responses (see Table 4 and the second column of Figure 2), the interaction effects are significant ($p's < .05$) but in the *opposite direction* in all but one of the vignettes (*No-Platforming*, $p = .87$).

Table 3. First set of responses: The effects of SDO, Condition, and their interaction on speech protection for the first sets of responses to each of the four vignettes in Study 2.

	<u>Online Shaming</u>			<u>Academic Freedom</u>			<u>Job Security</u>			<u>No-Platforming</u>		
	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		4.06	[3.92, 4.20]		4.72	[4.61, 4.83]		4.68	[4.50, 4.86]		4.03	[3.88, 4.19]
SDO	.02	.02	[-.07, .11]	-.03	-.02	[-.09, .06]	-.08	-.09	[-.21, .03]	.21***	.21	[.10, .31]
SpeechContent	.09	.11	[-.03, .24]	.02	.02	[-.09, .13]	-.15**	-.25	[-.43, -.07]	-.11	-.15	[-.31, .00]
SDO*SpeechContent	.30***	.25	[.16, .34]	.31***	.22	[.14, .29]	.27***	.29	[.17, .41]	.24***	.23	[.13, .34]

Note. SDO is a participant's score on the Social Dominance Orientation scale, mean centered. SpeechContent is coded 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

Table 4. Second set of responses: The effects of SDO, Condition, and their interaction on speech protection for the second sets of responses to each of the four vignettes in Study 2.

	<u>Online Shaming</u>			<u>Academic Freedom</u>			<u>Job Security</u>			<u>No-Platforming</u>		
	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		4.03	[3.91, 4.14]		4.74	[4.62, 4.86]		4.68	[4.51, 4.85]		4.10	[3.96, 4.24]
SDO	-.01	-.00	[-.08, .07]	-.04	-.03	[-.11, .05]	-.06	-.06	[-.18, .05]	.21***	.18	[.09, .27]
SpeechContent	-.14*	-.14	[-.26, -.03]	-.09	-.10	[-.02, .22]	.15*	.22	[.05, .39]	-.07	-.09	[-.23, .05]
SDO*SpeechContent	-.12*	-.08	[-.16, -.00]	-.16**	-.11	[-.19, -.04]	-.18**	-.18	[-.30, -.07]	-.00	-.01	[-.10, .09]

Note. SDO is a participant's score on the Social Dominance Orientation scale, mean centered. SpeechContent is coded 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

Table 5. Fixed effects from a mixed effects regression model predicting speech protection from SDO, Condition, Order, and their interactions across the four vignettes in Study 2. The model also includes random intercepts for each subject.

	<u>Online Shaming</u>			<u>Academic Freedom</u>			<u>Job Security</u>			<u>No-Platforming</u>		
	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		4.07	[3.92, 4.23]		4.72	[4.60, 4.84]		4.70	[4.52, 4.88]		4.07	[3.91, 4.23]
SDO	.01	.01	[-.09, .11]	.00	.00	[-.08, .08]	-.08	-.09	[-.21, .03]	.28***	.25	[.15, .36]
SpeechContent	.09	.11	[-.05, .26]	.01	.01	[-.10, .13]	-.16**	-.26	[-.44, -.08]	-.12*	-.17	[-.33, -.01]
Order	-.02	-.04	[-.17, .09]	.02	.05	[-.05, .14]	-.00	-.00	[-.11, .10]	.03	.09	[-.05, .24]
SDO*SpeechContent	.34***	.26	[.16, .36]	.37***	.25	[.17, .33]	.27***	.29	[.17, .41]	.26***	.23	[.13, .34]
SDO*Order	-.02	-.02	[-.10, .07]	-.01	-.01	[-.07, .05]	.02	.02	[-.04, .09]	-.02	-.02	[-.12, .08]
SpeechContent*Order	-.14	-.23	[-.51, .05]	-.06	-.10	[-.31, .12]	.22**	.49	[.15, .84]	.03	.06	[-.22, .35]
SDO*SpeechContent*Order	-.31***	-.34	[-.52, -.16]	-.39***	-.37	[-.51, -.23]	-.31***	-.47	[-.70, -.24]	-.18*	-.23	[-.42, -.04]

Note. Condition is coded 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. Order is coded 0 for the first time responding, 1 for the second. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

To investigate order effects more directly, we ran analyses with all responses included, and with an added variable called Order indicating whether a response is a participant's first or second response for that vignette (see Table 5).

The three-way interaction between SDO, SpeechContent and Order describes how the selective protection-effect (i.e. the interaction between SDO and SpeechContent on judgments) changes when a response is second in the order rather than first. Mirroring the findings from the analyses on the second set of responses only, there are significant three-way interactions for all the vignettes ($p < .001$, except for the *No-platforming vignette* which has $p = .016$).

Next, directly assessing participants' explicit support for content-sensitive speech norms, we considered their responses to being asked directly if the change in speech content made a difference to their judgments. The mean response across scenarios was 1.94 on a scale from 1 to 5 ($SD = 1.09$), meaning that participants largely reported that the changes in speech content were *not* morally relevant to them, indicating explicit support for content-neutral speech norms. Consistent with how the order effect was larger for the *No-platforming* scenario, the mean response here was significantly higher than for the three others, at 2.14 ($p < .01$).

In summary, when looking only at responses to the second, reversed set of scenarios, the effects of SDO on free speech judgments appear to be opposite of what one would expect if this was the first set. There is now increased protection of the speech that participants *oppose*, and decreased protection of the speech they *support*, demonstrating that participants comply with content-neutral speech norms when their practice of content-sensitive norms would otherwise be obvious. However, the ideologically consistent effects of SDO on judgements in the first set of scenarios again demonstrates that this is not the case if selective judgments may go undetected, even though participants explicitly deny that speech content matters for their endorsements of speech-rights.

Figure 2. Study 2. Interaction plots, showing the three-way interactions between SDO, speech content and presentation order on speech protection, across all four vignettes.

Note. SDO is Social Dominance Orientation. Speech protection is a participant's mean response to the four questions for each vignette about the blameworthiness of counteracting speech. There is one row for each vignette. The left column contains only responses from the first showing of a vignette, the right column contains only responses from the second showing. The darkness of shading around the regression lines represents the speech condition (bright gray is contra-SDO speech, dark gray is pro-SDO speech). The shaded areas around the regression lines indicate 95% confidence intervals. To accentuate which lines in the first- and second columns "consist of the same participants", the shapes of the regression lines indicate the between subject order manipulation: dotted lines are from the participants who got the Contra-SDO condition first, solid lines are those who got the Pro-SDO condition first.

4. Study 3

Study 3 was a pre-registered replication of Study 2. Additionally, we explored if the results of Study 2 were robust to the following set of potential moderators related to honesty and self-awareness: Introspectiveness (Hansell & Mechanic, 1985), the Honesty/Humility dimension of the HEXACO-PI (Ashton & Lee, 2009), and the number of heads reported out of five coin-flips when there is a bonus-payment for each head. For specific details on how we predicted these variables to moderate order effects, and for analyses investigating these predictions, see the Online Supplement.

4.1 Participants

To gain more accurate estimates of the effects in Study 2, we now increased the sample size from 300 to 400. Once again, an MTurk sample was specified to be 50-50 democrat and republican. Exclusion-criteria also mirrored those in Study 2. This yielded a mean sample size of 344 participants per analysis.

Among the 351 participants not fully excluded from the study ($M_{\text{age}} = 40.1$; $SD_{\text{age}} = 11.6$; %females = 42.1), there were 81% Caucasians, 4% Asians, 6% African Americans, 6% Hispanics, and 3% other.

4.2 Procedure

The procedure from Study 2 was repeated, except that now the Introspectiveness- and the Honesty/Humility scales were filled out following the SDO scale. Additionally, after the subjects had responded to all the vignettes, they were asked to flip a coin five times and report the number of heads. They were informed that they would receive 10 cents for each head they reported.

4.3 Results

For all four vignettes, the two-way interactions between SDO and Condition and also the three-way interactions between SDO, Condition, and Order were all significant at $p < .0001$, in the predicted directions (see Table 6 and Figure 4). The effect sizes were mostly similar to those in Study 2. Neither introspection, trait-honesty, nor cheating for an economic bonus were found to moderate these effects. This suggests either that the processes producing these biases operate too far outside consciousness to be accessible through introspection, or that participants are uninterested in overriding them when they gain access to them (regardless of their levels of honesty/integrity). For further analyses and discussion of Study 3, see the online supplement.

Table 6. Fixed effects from a mixed effects regression model predicting speech protection from SDO, Condition, Order, and their interactions across the four vignettes in Study 3. The model also includes random intercepts for each subject

	<u>Online Shaming</u>			<u>Academic Freedom</u>			<u>Job Security</u>			<u>No-Platforming</u>		
	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		4.07	[3.93, 4.20]		3.52	[3.40, 3.63]		4.32	[4.17, 4.47]		3.65	[3.52, 4.26]
SDO	-.04	-.03	[-.12, .06]	.04	.03	[-.05, -.02]	-.01	-.01	[-.11, .10]	-.18***	.16	[.07, .26]
SpeechContent	.01	.01	[-.12, .15]	.06	.07	[-.19, .23]	-.15**	-.22	[-.37, -.07]	-.17**	-.22	[-.37, -.09]
Order	.03	.06	[-.06, .18]	.01	-.02	[-.10, .30]	-.00	-.00	[-.08, .08]	.06*	.16	[.03, .28]
SDO*SpeechContent	.31***	.24	[.16, .33]	.34***	.27	[.18, .35]	.27***	.28	[.18, .39]	.35***	.32	[.22, .41]
SDO*Order	.02	.02	[-.06, .10]	-.01	-.01	[-.07, .04]	-.02	-.03	[-.08, .03]	-.04	-.05	[-.13, .04]
SpeechContent*Order	-.08	-.14	[-.38, .10]	-.08	.13	[-.10, .35]	.17*	.36	[.07, .65]	.07	.13	[-.11, .37]
SDO*SpeechContent*Order	-.26***	-.29	[-.45, -.13]	-.34***	-.38	[-.53, -.22]	-.32***	-.48	[-.68, -.28]	-.26***	-.33	[-.50, -.17]

Note. Condition is coded 1 for the Pro-SDO condition and -1 for the Contra-SDO condition. Order is coded 0 for the first time responding, 1 for the second. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

Figure 3. Study 3. Interaction plots, showing the three-way interactions between SDO, speech content and presentation order on speech protection, across all four vignettes.

Note. SDO is Social Dominance Orientation. Speech protection is a participant's mean response to the four questions for each vignette about the blameworthiness of counteracting speech. There is one row for each vignette. The left column contains only responses from the first showing of a vignette, the right column contains only responses from the second showing. The darkness of shading around the regression lines represents the speech condition (bright gray is contra-SDO speech, dark gray is pro-SDO speech). The shaded areas around the regression lines indicate 95% confidence intervals. To accentuate which lines in the first- and second columns "consist of the same participants", the shapes of the regression lines indicate the between subject order manipulation: dotted lines are from the participants who got the Contra-SDO condition first, solid lines are those who got the Pro-SDO condition first.

5. Study 4

Up to this point, preferences for speech norms have only been assessed in relation to specific scenarios, rather than in general terms. Evaluating abstract moral principles can be challenging; Having them applied to specific scenarios makes things more concrete and understandable, and it might give participants better access to their attitudes. Still, we feel it is worth investigating if the expressed support for content-neutral speech norms in specific scenarios carries over to support for the more general principle that the content of speech should not affect judgments of its suppression.

Furthermore, we investigate the prediction that participants are particularly supportive of content-neutral speech norms when proposed by the ideological opposition. After all, the benefits to oneself from having others give equal protections to all speech are largest when these others would have otherwise gone after views that you support. To explore this, we ask participants to take sides in a hypothetical discussion about free speech between two friends, one pro- and one against content-sensitive speech norms (see Box 2 for the full text). We experimentally manipulate the ideological commitments of *both* discussants simultaneously, so that they are either both liberals or both conservatives. We predict that the level of support for content-neutral speech norms will be higher, and that support for content-sensitive speech norms will be lower, when they are proposed by someone who does not share the participant's own ideological commitments.

To address whether participants are open about supporting content-neutral norms more when they are voiced by the opposition, we again implement an order manipulation; Participants respond to both conditions of the vignette, in randomized order. This also addresses the possibility that participants might be stating their support for the particular brand of content-sensitivity that comes from having a specific political outlook, rather than them stating support for content-sensitivity in general (as we intend them to).

5.1 Participants.

Even though the sample size of <300 was sufficient for achieving high power to detect the predicted effects in Study 2, we aimed at recruiting 400 participants for Study 4 because we effect sizes might be attenuated by our design modifications to probe motivated general, abstract endorsement of free speech. Once again, we recruited from Amazon's Mechanical Turk a sample that was split 50-50 between republicans and democrats, using MTurk's filtering functionality.

Out of the 406 people who completed our survey, 33 failed the attention check, and one did not respond to our predictor variable, leaving a sample of 372 participants for our analyses ($M_{\text{age}} = 39.1$, $SD_{\text{age}} = 12.2$, 213 females). This sample consisted of 82% Caucasians, 9% African Americans, 4% Asians, and 3% Hispanics.

5.2 Procedure.

Participants gave an informed consent and basic demographic information, and they indicated their political position on economic- and social issues using two visual analog scales ranging from “very liberal” to “very conservative”. The scores on these two scales ($r = .81$) were combined to form an overall index of political orientation functioning as our predictor variable. Participants then read and responded to the vignette.

The names of the two discussants in the vignette, Riley and Casey, were chosen to be as uninformative as possible with regards to both gender and political affiliation: People named Riley or Casey in the US are about equally likely to be male or female (Flowers, 2015), and they are also about equally likely to have voted democratic or republican in the 2012 election (Clarity Campaign Labs, 2018).

The responses to the questions were combined to create a single scale indicating support for content-neutral speech norms ($\alpha = .91$). Higher scores on this scale then imply agreement with Casey, who advocated content-neutral norms, and disagreement with Riley, who took the opposing position.

After responding to the vignette the first time, participants were asked to read the vignette once again, this time in the other experimental condition, and they once again rated their agreement to the same six statements about the vignette. They were also asked to explicitly indicate the extent to which the changes to the text would impact their responses.

Box 2. Vignette in Study 4

Riley and Casey, who are both politically **[conservative/liberal]**, were discussing free speech issues. They debated various actions that have been taken against people spreading their views. Such as sabotaging speaking events, getting people fired from their jobs, or barraging people with shame and insults on social media.

Riley argued that the wrongfulness of such actions depends on which particular views are being targeted. Certain views are more damaging to society than others, Riley argued, so therefore it is more morally justified to use harsh measures to counteract the spreading of these views. As examples of positions that are damaging to society, Riley used statements made by **[left/right]** wing pundits and -activists and their sympathizers.

1. Riley is right that the legitimacy of using harsh measures to counter a view depends on how harmful the view is to society. [Reversed]
2. Riley's position is hypocritical.
3. Riley's position is stifling to the free exchange of ideas.

Casey rather argued that actions to counteract the spreading of views are equally right or wrong regardless of which particular view is being spread. While being strongly **[pro-gun/anti-gun]** in the gun control debate, Casey still argued that sabotaging an **[anti-gun/pro-gun]** speaking event is morally equivalent to sabotaging a **[pro-gun/anti-gun]** event. Either both actions are wrong or both are right, but it can't be right to sabotage one but not the other, Casey claimed.

1. Casey is correct that the wrongfulness of silencing other people's views should not depend on what those views are.
2. Casey's position leaves society too open to the spreading of toxic ideologies. [Reversed]
3. Casey is going too far in a misunderstood attempt at seeming fair and consistent. [Reversed]

Note. Full vignette used in Study 3. Bold font and brackets indicate points were the two experimental conditions differ, with the text for the Conservative Condition first, and then the text for the Liberal Condition after the slash. The numbered points are the statements participants rated their agreements with, on 7-point Likert scales.

5.3 Results.

Consistent with studies 1, 2, and 3, our sample sided quite clearly with the discussant arguing for content-neutral speech norms, with a mean agreement here on the first reading of the vignette at 5.1 on a scale from 1 to 7. This was significantly higher than the scale midpoint of 4.0 ($t(371) = 14.8$, $p < .001$).

In our model predicting support for content-neutral speech norms (see Table 7 and Figure 4), the coefficient for the interaction was significantly negative ($p < .001$), suggesting that left-leaning participants were relatively more inclined to support such norms when both discussants are conservatives, and that right-leaning participants had the opposite tendency. The positive main effect of political orientation ($p < .001$) suggests that, averaging across the conditions, conservative participants were more strongly supportive of content-neutral norms than liberals were.

In the models investigating order effects (Table 8) we find, as in studies 2 and 3, that the coefficient for the three-way interaction is of opposite sign and of larger magnitude than the coefficient for the interaction between political orientation and Condition ($p < .001$), indicating that order of presentation has a substantial effect on responses here as well.

In summary, while participants generally and clearly favor content-neutral norms (conservatives in our sample more so than liberals), their levels of support nevertheless systematically depend on which political coalition was espousing them. Specifically, they support content-neutral speech norms *more* when they are proposed by someone with the *opposite* political affiliation as their own. This conforms to how content-neutral norms are particularly useful when supported by those who would otherwise selectively protect speech you oppose and restrict speech you support. However, this selective endorsement of content-neutral speech norms is again significantly attenuated when this selectivity would be laid bare by answers to the second, ideologically-reversed scenario, again indicating that people are motivated to appear morally consistent.

Table 7. The effects of Political Orientation, Condition, and their interaction on levels of expressed support for content-neutral speech norms.

	<i>Support for content-neutral norms</i>		
	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		5.18	[5.05, 5.32]
Pol_Orient	.22***	.09	[.05, .13]
Condition	.03	.05	[-.09, .19]
Pol_Orient*Condition	-.21***	-.09	[-.13, -.05]

Note. Pol_Orient is political affiliation from -5 (“very liberal”) to 5 (“very conservative”). Condition is coded 1 for the Conservative condition, where both discussants are conservatives, and -1 for the Liberal condition, where they are both liberals. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$.

Table 8. Fixed effects from a mixed effects regression models predicting expressed support for content-sensitive and –neutral speech norms from political orientation, Condition, Order, and their interactions. The model also includes random intercepts for each participant.

	<i>Support for content-neutral norms</i>		
	β	<i>b</i>	<i>CI</i> _{95%}
(Intercept)		5.18	[5.05, 5.32]
Pol_Orient	.23***	.09	[.05, .13]
Condition	.03	.05	[-.09, .19]
Order	.02	.06	[-.01, .12]
Pol_Orient*Condition	-.22***	-.09	[-.13, -.05]
Pol_Orient *Order	-.04*	-.02	[-.04, -.00]
Condition*Order	-.05	-.10	[-.37, .16]
Pol_Orient*Condition*Order	.23***	.13	[.05, .21]

Note. Pol_Orient is political orientation measured from -5 (“very liberal”) to 5 (“very conservative”). Condition is coded 1 for the Conservative condition, where both discussants are conservatives, and -1 for the Liberal condition, where both are liberals. Confidence intervals refer to the unstandardized betas. * $p \leq .05$, ** $p \leq .01$, *** $p \leq .001$. Degrees of freedom for significance tests are calculated using the Kenward-Roger approximation.

Figure 4. Combined scatter- and interaction plots, showing the three-way interactions between Political Orientation, Condition and Order condition on Level of Agreement, across all four vignettes.

Note. Political Orientation is a participant's position from very liberal (-50) to very conservative (50). Level of Agreement is a participant's mean response to the six statements indexing support for content-neutral speech norms. The left column contains only responses from the first showing of a vignette, the right column contains only responses from the second showing. The shade around regression lines represents the condition (bright gray is when discussants are liberals (Liberal Condition), dark gray is when they are conservatives (Conservative Condition)). To accentuate which lines in the first- and second columns consist of the same participants, the shapes of the regression lines indicate the between subject order manipulation: dotted lines are from the participants who got the liberal condition first ("LibFirst"), solid lines are those who got the Conservative Condition first ("ConFirst").

6. Discussion

The present series of studies qualifies, and consolidates, emerging findings that people are selectively protective of views they support when making free speech judgments. Among the competing hypotheses that people either (A) operate with explicitly different standards for different speech (*content-sensitive speech norms*), or (B) seek to maintain a façade of operating with a single standard for all speech (*content-neutral speech norms*) while selectively adjusting this standard to suit the context, the present investigation lends support to (B). This fits well with the more general notion that people tactically adjust their moral judgments (Bartels, 2008), as well as the notion that people lack awareness of the flexibility of their moral convictions and rather feel like steadfast believers in the values and norms that happen to be currently helpful (Trivers, 2000; Ditto et al., 2018).

In Study 1 we showed that when judging a series of scenarios raising free speech issues, the tendency to selectively protect supported speech was dampened when participants had been previously exposed to other scenarios with speech of the opposite ideological flavor, suggesting that participants want to seem as if they employ the same standard for all speech.

Studies 2 and 3 further corroborated this finding by investigating *order effects*: when participants responded to both conditions of our design sequentially, without knowing that the second condition was coming when answering the first, and without being able to go back to alter prior responses, the order of presentation was found to have a large impact on responses. To seem consistent, participants gave similar responses to the second version of each scenario as those they gave for the first. And since their first responses were biased in favor of speech they supported, their second responses then produced almost the same level of bias in favor of speech they opposed.

Study 4 provided additional nuance, by investigating if people are selective in their approval of content-neutral speech norms, depending on the political ideology of the person proposing them. By manipulating the ideological affiliation of the discussants in a hypothetical discussion about speech norms, we found that the tendency to support the content-neutral position, while strong across both conditions, was substantially intensified when discussants occupied the opposite side of the liberal-conservative spectrum to one's own. Furthermore, we find that people are motivated to hide this selectivity as well, as we once again observed an order effect.

If our results are generalized to other domains of moral and social judgment, they suggest that it is normal to have at least two sets of norms and values: One set is the “official” one, which we present to others; the other set reflects our true preferences, which we keep hidden (perhaps even from ourselves). Our publicly endorsed norms and values will typically represent a compromise between the conflicting interests of many parties. For example, content-neutral speech norms can be seen to represent a compromise between speech norms that are explicitly content-sensitive in either one or the other ideological direction. These kinds of norms can be seen as our offer in a negotiation with parties who have opposing interests to our own; if these parties are willing to adjust, then we are willing to

return the favor, and settle on a set of rules everyone can live with. When there are situations giving room for plausible deniability, however, such that we no longer need to compromise, and can instead safely work towards our true goals, we take advantage of this. This way, we are at the same time able to hold the ideological opposition accountable whenever they detectably transgress on shared norms, since we can claim that we are keeping our end of the bargain, yet we are also able to transgress on these very norms ourselves whenever we can do so undetectably.

A potential ambiguity with our series of experiments concerns whether participants were truly self-deceived or if they rather were to some extent conscious about the selective nature of their judgments. In principle, it is possible that participants at all times had conscious access to their “true judgments” about the issues in our scenarios, and that they purposely adjusted away from these true judgments to suit the context. But this kind of deliberate deceit seems implausible to us, since the survey was completely anonymized and participants stood to gain nothing from such impression management. We see it as more likely that most of our participants did believe that speech content was irrelevant to their judgments, even as they unconsciously took it into account. Such a conceptualization of moral judgment processes as being largely hidden from our conscious minds has broadening support (Haidt, 2012; Mercier & Sperber, 2011; Trivers, 2011).

All the views expressed in our hypothetical scenarios are currently quite well-supported in the US where our data was collected. This was deliberate, as we needed substantial amounts of participants to be both for and against all the different views presented in order for our analyses to be informative. A drawback to this, however, is that all of the views in our scenarios can be argued to be inside the so-called “Overton Window”, containing all the views within the realm of acceptable discourse (Lehman, 2012). Therefore, we cannot generalize our findings outside of this window, to fringe views widely seen as truly toxic and destructive. These kinds of views are hard to study with our design, since it is hard to find samples that include a sufficient amount of supporters, and since review boards will tend to suggest their removal from surveys. Our intuitive prediction is that order effects would be attenuated, or perhaps non-existent, if our design included views most people find truly repulsive: While many support equal speech-rights for a broad range of views, this range might not be infinitely broad. Beyond this range, harsher measures could become more acceptable. Notably, the scenario producing the weakest order effect in Study 2 - the public speaker who is either for or against multiculturalism – perhaps raises the most polarizing current topic. Interestingly, both sides see the views of the other side to be about equally far beyond the pale here, suggesting that rather than discussing a single Overton Window for the US as a whole, it is perhaps more accurate to conceptualize multiple, partly overlapping windows for separate sub-groups.

Following from this, it can be argued that as divisions between mainstream factions increase, as is arguably the case in the US in recent years (Garimella & Weber, 2017), explicitly content-sensitive speech norms in support of coalition-specific *sacred values* (Tetlock, 2003) might rise in popularity among all camps as the hopes of sustaining cooperation across ideological divides

decrease. An optimistic take-away from our results is then that among American participants polarization seems to not (yet?) have reached the point where most people have completely given up on having a shared set of norms.

7. Conclusion

While judgments of actions to restrict speech systematically depend on whether one supports its ideological content, we show that people seek to appear as if this is not the case for them. When this appearance of impartiality is threatened, people adjust their judgments and extend similar protections to speech they support and oppose.

To the extent that people truly wish to follow this precept of not taking speech content into account when making judgments about speech rights, our study suggests that it is not sufficient to just go with one's intuitions. The processes producing our intuitions seem to be more tactical and devious than we are able to tell from introspection. A good approach could then be to imagine oneself as blind to the specifics of each case; how would one feel about speech being counteracted in a certain way if the content of this speech was unknown? Recent evidence suggests that this kind of veil-of-ignorance reasoning can indeed promote moral impartiality (Huang, Greene, & Bazerman, 2019).

References

- Alexander, S. (2014a). In Favor of Niceness, Community, and Civilization. Retrieved from <http://slatestarcodex.com/2014/02/23/in-favor-of-niceness-community-and-civilization/>
- Alexander, S. (2014b). Beware Isolated Demands For Rigor. Retrieved from <http://slatestarcodex.com/2014/08/14/beware-isolated-demands-for-rigor/>
- Alexander, S. (2018). Conflict vs. mistake. Retrieved from <https://slatestarcodex.com/2018/01/24/conflict-vs-mistake/>
- Ashton, M. C., & Lee, K. (2009). The HEXACO–60: A short measure of the major dimensions of personality. *Journal of personality assessment*, 91(4), 340-345.
- Bartels, D. M. (2008). Principled moral sentiment and the flexibility of moral judgment and decision making. *Cognition*, 108(2), 381-417.
- Bey, A. (2017). Does Your Right to Free Speech Extend to the Workplace?. Retrived from <https://www.mediabistro.com/be-inspired/right-free-speech-workplace/>
- Blumer, H. (1958). Race prejudice as a sense of group position. *Pacific sociological review*, 1(1), 3-7.
- Brandt, M. J., Chambers, J. R., Crawford, J. T., Wetherell, G., & Reyna, C. (2015). Bounded openness: The effect of openness to experience on intolerance is moderated by target group conventionality. *Journal of Personality and Social Psychology*, 109(3), 549.
- Brennan, O. (2015). The Enemy Control Ray. Retrieved from <https://thingofthings.wordpress.com/2015/05/15/the-enemy-control-ray/>
- Campbell, B. (2018). The Free Speech Crisis on Campus Is Worse than People Think. Retrived from <https://quillette.com/2018/11/14/the-free-speech-crisis-on-campus-is-worse-than-people-think/>
- Clarity Campaign Labs (2018). How Democratic is Your Name? Retrieved from <http://www.claritycampaigns.com/names/>
- Crawford, J. T. (2014). Ideological symmetries and asymmetries in political intolerance and prejudice toward political activist groups. *Journal of Experimental Social Psychology*, 55, 284-298.
- Crawford, J. T., & Pilanski, J. M. (2014). Political intolerance, right and left. *Political Psychology*, 35(6), 841-851.
- Crawford, J. T., & Xhambazi, E. (2015). Predicting political biases against the Occupy Wall Street and Tea Party movements. *Political Psychology*, 36(1), 111-121.
- Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, 1(11), 769.
- Crook, C. (2017) Google moves into the business of thought control. Retrieved from <https://www.bloomberg.com/opinion/articles/2017-08-14/google-moves-into-the-business-of-thought-control>
- Disinvitation Database - FIRE. (2019). Retrieved from <https://www.thefire.org/resources/disinvitation-database/>

- Ditto, P. H., Liu, B. S., Clark, C. J., Wojcik, S. P., Chen, E. E., Grady, R. H., ... & Zinger, J. F. (2018). At least bias is bipartisan: A meta-analytic comparison of partisan bias in liberals and conservatives. *Perspectives on Psychological Science*, 1745691617746796.
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, 39(2), 175-191.
- Flowers, A. (2015). The Most Common Unisex Names In America: Is Yours One Of Them? Retrieved from <https://fivethirtyeight.com/features/there-are-922-unisex-names-in-america-is-yours-one-of-them/>
- Galef, J. (2018). Insightful articles on free speech & social justice. Retrieved from <https://juliagalef.com/2018/01/10/insightful-articles-on-free-speech-social-justice/>
- Garimella, V. R. K., & Weber, I. (2017). A long-term analysis of polarization on Twitter. arXiv preprint arXiv:1703.02769.
- Graham, P. (2004). What you can't say. Retrieved from <http://www.paulgraham.com/say.html>
- Green, J., & Karolides, N. J. (2014). *Encyclopedia of censorship*. Infobase Publishing.
- Greene, J. D. (2014). *Moral tribes: Emotion, reason, and the gap between us and them*. Penguin.
- Haidt, J. (2012). *The righteous mind: Why good people are divided by politics and religion*. Vintage.
- Hansell, S., & Mechanic, D. (1985). Introspectiveness and adolescent symptom reporting. *Journal of Human Stress*, 11(4), 165-176.
- Ho, A. K., Sidanius, J., Pratto, F., Levin, S., Thomsen, L., Kteily, N., & Sheehy-Skeffington, J. (2012). Social dominance orientation: Revisiting the structure and function of a variable predicting social and political attitudes. *Personality and Social Psychology Bulletin*, 38(5), 583-606.
- Ho, A. K., Sidanius, J., Kteily, N., Sheehy-Skeffington, J., Pratto, F., Henkel, K. E., ... & Stewart, A. L. (2015). The nature of social dominance orientation: Theorizing and measuring preferences for intergroup inequality using the new SDO₇ scale. *Journal of Personality and Social Psychology*, 109(6), 1003.
- Huang, K., Greene, J. D., & Bazerman, M. (2019). Veil-of-ignorance reasoning favors the greater good. *Proceedings of the National Academy of Sciences*, 116(48), 23989-23995.
- Kahan, D. M. (2015). The politically motivated reasoning paradigm, part 1: What politically motivated reasoning is and how to measure it. *Emerging trends in the social and behavioral sciences: An interdisciplinary, searchable, and linkable resource*, 1-16.
- Kant, I. (1797/2013). *Moral law: Groundwork of the metaphysics of morals*. Routledge.
- Kenward, M. G., & Roger, J. H. (1997). Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, 983-997.

- Kleppestø, T. H., Czajkowski, N. O., Vassend, O., Røysamb, E., Eftedal, N. H., Sheehy-Skeffington, J., ... & Thomsen, L. (2019). Correlations between social dominance orientation and political attitudes reflect common genetic underpinnings. *Proceedings of the National Academy of Sciences*, 116(36), 17741-17746.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological bulletin*, 108(3), 480.
- Lehman, J. (2012). A Brief Explanation of the Overton Window. Retrieved from <http://www.mackinac.org/OvertonWindow>
- LeVine, R. A., & Campbell, D. T. (1972). Ethnocentrism: Theories of conflict, ethnic attitudes, and group behavior.
- Lindner, N. M., & Nosek, B. A. (2009). Alienable Speech: Ideological Variations in the Application of Free-Speech Principles. *Political Psychology*, 30(1), 67-92.
- 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(5), 665.
- McClelland, J. (1996). The place of Elias Canetti's Crowds and power in the history of Western social and political thought. *Thesis eleven*, 45(1), 16-27.
- Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and brain sciences*, 34(2), 57-74.
- Mill, J. S. (1966). On liberty. In *A selection of his works* (pp. 1-147). Palgrave, London.
- Munroe, R. (2014). Free Speech. Retrieved from <https://xkcd.com/1357/>
- Petersen, M. B. (2015). Evolutionary political psychology: On the origin and structure of heuristics and biases in politics. *Political Psychology*, 36, 45-78.
- Pinker, S. (2007). In defense of dangerous ideas. *Chicago Sun-Times*, 15.
- Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of personality and social psychology*, 67(4), 741.
- R Core Team (2017). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Sidanius, J. (1993). The psychology of group conflict and the dynamics of oppression: A social dominance perspective. In S. Iyengar & W. J. McGuire (Eds.), *Duke studies in political psychology. Explorations in political psychology* (pp. 183-219). Durham, NC, US: Duke University Press.
- Sidanius, J., & Pratto, F. (2004). Social dominance theory: A new synthesis.
- Simler, K., & Hanson, R. (2017). *The Elephant in the Brain: Hidden Motives in Everyday Life*. Oxford University Press.

- Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in cognitive sciences*, 7(7), 320-324.
- Tetlock, P. E. (2007). Psychology and politics: The challenges of integrating levels of analysis in social science.
- Thomas, A. J., Thomsen, L., Lukowski, A. F., Abramyan, M., & Sarnecka, B. W. (2018). Toddlers prefer those who win, but not when they win by force. *Nature Human Behavior Books (AZ)*.
- Thomsen, L., Frankenhuys, W. E., Ingold-Smith, M., & Carey, S. (2011). Big and mighty: Preverbal infants mentally represent social dominance. *science*, 331(6016), 477-480.
- Thomsen, L., Green, E. G., Ho, A. K., Levin, S., van Laar, C., Sinclair, S., & Sidanius, J. (2010). Wolves in sheep's clothing: SDO asymmetrically predicts perceived ethnic victimization among White and Latino students across three years. *Personality and Social Psychology Bulletin*, 36(2), 225-238.
- Tooby, J., & Cosmides, L. (2010). Groups in mind: The coalitional roots of war and morality. *Human morality and sociality: Evolutionary and comparative perspectives*, 91-234.
- Trivers, R. (2000). The elements of a scientific theory of self-deception. *Annals of the New York Academy of Sciences*, 907(1), 114-131.
- Trivers, R. (2011). The folly of fools: The logic of deceit and self-deception in human life. Basic
- Von Hippel, W., & Trivers, R. (2011). Reflections on self-deception. *Behavioral and Brain Sciences*, 34(1), 41-56.
- Weeden, J., & Kurzban, R. (2014). *The hidden agenda of the political mind: How self-interest shapes our opinions and why we won't admit it*. Princeton University Press.
- White II, M. H., & Crandall, C. S. (2017). Freedom of racist speech: Ego and expressive threats. *Journal of personality and social psychology*, 113(3), 413.