How Polarized Are Online Conversations About Childhood?

Breanna E. Green[1], **William R. Hobbs**[2]

[1] Cornell University, Department of Information Science, Ithaca, NY, USA
[2] Cornell University, Departments of Government and Psychology, Ithaca, NY, USA
BEGreen@infosci.cornell.edu, Hobbs@cornell.edu

Abstract

2020 through 2023 were unusually tumultuous years for children in the United States, and children's welfare was prominent in political debate. Theories in moral psychology suggest that political parties would treat concerns for children using different moral frames, and that moral conflict might drive substantial polarization in discussions about children. However, such partisan frames may still differ very little if there is limited underlying disagreement about moral issues and everyday concerns in childhood when not explicitly referencing politics. We evaluate claims of universality and division in moral language using tweets from 2019-2023 linked to U.S. voter records, focusing on *expressed* morality. Our results show that mentions of children by Republicans and Democrats are usually similar, differing no more than mentions by women and men, and tend to contain no large differences in accompanying moral words. To the extent that mentions of children did differ across parties, these differences were constrained to topics polarized well before the pandemic – and slightly heightened when co-mentioned with 'kids' or 'children'. These topics reflected a small fraction of conversations about children. Overall, polarization of online discussion around childhood appears to reflect escalated polarization on lines of existing partisan conflicts rather than concerns originating from new concerns about the welfare of children during and after the pandemic.

Code for this paper — https://osf.io/j93az/

Introduction

Have online conversations about children become far more polarized since 2019? If so, was increased polarization driven by new concerns about the well-being of children during the pandemic? Or was polarization more narrowly limited to topics that were deeply partisan long before 2020?

In the United States, children's welfare was a leading concern during the COVID-19 pandemic, especially in terms of health (e.g., vaccinations) and negative academic impacts such as severe learning loss (Donnelly and Patrinos 2021; Skar, Graham, and Huebner 2022). In 2021 and 2022, Democrats and Republicans engaged in a number of politically charged debates about protecting children that transcended direct pandemic-related concerns, including the type of information that should be taught in schools (Pew Research Center 2022a). Further, the political rhetoric surrounding parental rights extended to current hotbed issues such as abortion, gender-affirming care, and critical race theory (CRT) (Cahn 2022; Pew Research Center 2022b). Numerous state and federal level proposed policies either reflected or drove such debates – including an anti-CRT campaign by Glenn Youngkin in Virginia and Floridian Governor Ron DeSantis's Stop the Wrongs to Our Kids and Employees (W.O.K.E.) Act, as well as the Democratic party's Invest in Child Safety and Raise the Age Acts.

Any political debate is likely to reflect moral values in one way or another (Ryan 2014), but debates about children might contain especially prominent moral frames. Yet, conversations about children are likely to be more *universally* focused on childhood innocence and the need to protect children (Haslam, Rothschild, and Ernst 2000; Woodrow 1999). Everyday parenting concerns often discussed online, such as childcare, screen time, and eating behaviors (Thornton et al. 2023) might be only rarely polarized.

Here, we evaluate the juxtaposition of everyday language and potentially universal moral concerns about protecting children with more divisive partisan frames on social issues. We focus on measuring the extent of polarization in online conversations referencing children. Specifically, we contextualize differences in partisan language with gendered differences and with co-mentions of political topics that were polarized long before the COVID-19 pandemic and later political debates. For this purpose, we use social media posts from 2019 through 2023 by Twitter (later renamed to '*X*') users whose profiles were linked to voter files that include their partisanship and demographics (Grinberg et al. 2019; Hughes et al. 2021; Shugars et al. 2021).

In this, we evaluate a claim that we suspect to be true, even during and after the pandemic: that Democrats and Republicans talk about children in the same ways for the vast majority of topics. Clearly illustrating this is nonetheless important. Core similarities might be overlooked (Hartman et al. 2022) due to strategic use of children's health and welfare in political rhetoric (e.g., to paint an opposition group as more extreme than they really are) and resulting misperceptions of out-party members (Wilson, Parker, and Feinberg 2020). Or due to a general tendency to misattribute disagreements (Ren and Schaumberg 2024). Because of this, highlighting commonalities might itself de-polarize some conversations (Ahler and Sood 2018; Hartman et al. 2022; Syropoulos and Leidner 2023; Levendusky 2023). Relatedly, research focused only on finding differences between parties – without work that can also demonstrate levels of similarities – may tend to contribute to such out-party misperceptions. Further, illustrating enduring political differences when they occur can also provide a justification for *acknowledging* (political) differences in conversation when attempting to persuade on politically adjacent topics (Kalla and Broockman 2020; Hartman et al. 2022; Ecker et al. 2022). This is important in cases where such interpersonal persuasion might improve children's welfare without needing to alter most political identities and beliefs.

In our analyses, we evaluate levels of polarization overall and over time using a variety of approaches: moral dictionaries, supervised feature extraction, and (distances in) word embeddings. A major focus in all analyses is not merely *finding* differences but in contextualizing the size of differences. Although we find it possible to mine text data for partisan differences in language use about children, most conversations about children by Democrats and Republicans online are similar. When non-politician Democrats and Republicans do differ slightly, differences seem to more consistently arise out of the expansion of pre-existing political debates – rather than, as seen during the pandemic, more novel concerns about the welfare of children specifically.

Background

Below, we provide background on partisan variation in moral values and language, and discuss how we might expect the setting of *any* conversation about children online[1] might differ from more focused survey items on political surveys

(e.g., Barker and Tinnick 2006) or from the rhetoric of politicians (Kraft and Klemmensen 2024). Drawing on past research, we consider whether to (1) expect new concerns about children during the pandemic to drive political division – perhaps through competing moral frames applied to all potentially political concerns – or (2) primarily reflect expanded conflict on already polarized issues such as racism, gender identity, immigration, and gun control.

Finally, we note conversations about children may differ by time and place. In this, we explain why we might expect Florida to be a most likely case for periods of elevated polarization in conversations about children.

Why Online Conversations About Children Might *Not* Polarize

Morality, Politics, and Children. Theories about morality in politics are useful because they help us abstract away from specific debates and political contexts, linking lines of conflict to more constant differences across political divides. For example, a debate about euthanasia or marijuana can be seen as a manifestation of recurring conflict about purity and sanctity of the human body, and tied to politics through religiosity and conservative ideology (Silver 2020).

[1]Here, specifically on Twitter.

In studying conversations about children, research on morality and politics can be especially informative – parenting, family, and views toward children play an important role in theories of moral psychology (e.g., Lakoff 1996; Haidt and Joseph 2004). Yet, past work might both suggest far-reaching universality in underlying concern for the well-being of children and also extensive polarization in framing and language about how best to protect and provide for children.

An especially prominent theory commonly used in analyses of moral language and politics called moral foundations theory (MFT; Haidt and Joseph 2004; Haidt and Graham 2007) argues liberals and conservatives understand morals and base values on five foundations: care/harm, fairness/reciprocity, (in-group) loyalty, authority/respect, and sanctity/purity. Research using MFT has found liberal morality to be strongly connected to the care/harm and fairness/reciprocity foundations, while conservative morality is loosely tied to all five foundations but more strongly connected to (in-group) loyalty, authority/respect, and sanctity/purity (Graham, Haidt, and Nosek 2009).

In this, the care/harm foundation is *itself* motivated through a connection to parenthood and hypothesized evolutionary motives to protect vulnerable others, especially children. In text analysis, terms 'child' and 'children' are regularly included as influential measures of care-related language. Consequently, and despite partisan differences, this theory and related methods based on it might suggest most moral concerns (and, by extension, political concerns) about children would tend to relate to care and protection of them (Haidt, Joseph et al. 2007). Thus potentially but not necessarily limiting the amount of division in language about child-related issues.

Online Social Media and Political Conversations. Most prior work on morality in language has focused on more explicitly moral scenarios and contexts. Applications of moral foundations theory to language often focus on politicians' speeches, debates, and social media accounts (Deason and Gonzales 2012; Reiter-Haas, Kopeinik, and Lex 2021; Hackenburg, Brady, and Tsakiris 2023; Brisbane, Hua, and Jamieson 2023). A similarly large body of work has studied more typical language related to moral dilemmas in everyday life (Kennedy et al. 2021; Nguyen et al. 2022; Atari et al. 2023) and specific public issues (e.g., vaccines, gun control, same-sex marriage, and climate change) (Weinzierl and Harabagiu 2022; Brady et al. 2017a).

A notable advantage of using data from online social media is that we might expect our findings to differ from previous work due to our focus on very general language. Though, perhaps surprisingly, Pew Research Center (2022c) estimates 1/3 of U.S. adults tweets are political, the majority of online conversations are *not* about politics. Therefore, we might expect this children *and* explicit politics subset to be relatively small. Conversations might be far more likely to cover, for example, childcare, screen time, and meals (Thornton et al. 2023). And so even if there is substantial polarization in political conversations, then their influence on overall levels of polarization in online conversations about children may be limited.

In considering online social media and moral or political conversations, we should note there is an impressive body of machine learning research on classification of moral statements in texts (e.g., Garten et al. 2016; Lin et al. 2018; Huang, Wormley, and Cohen 2022). Our goals here differ from that line of work, as we study how two groups (as linked to voter file demographics) differ *on average*. Precision in our estimates will be largely derived from sample size rather than highly accurate predictions at the document level.

Why Online Conversations About Children Might (*Sometimes*) Polarize

The primary reason we might expect polarization in online conversations is from news reports about the effects of the pandemic, responses to pandemic policies, and waves of new state-level policies related to children and culture war issues.[2] There has been seemingly extensive coverage of masking policies, CRT, vaccines, gender identity, remote schooling, and unusually fractious school board meetings. However, such coverage is unlikely to provide a clear answer about the extent of polarization in everyday online conversations. Media coverage is often a poor indicator of the underlying prevalence of a phenomenon (Boydstun 2013), and news reports tend to be focused on particularly novel and news-worthy information. They avoid covering mundane everyday activities, which may or may not be affected by political trends and current events.

Beyond reporting and anecdotes, a number of theories about politics and political polarization have long argued that parenting views reflect, and perhaps drive, political attitudes (e.g., Lakoff 1996). There is some empirical work supporting the claim that endorsement of disciplinarian versus nurturing *parenting* styles is related to partisanship as well as a broad range of conservative and liberal political preferences (Barker and Tinnick 2006; Deason and Gonzales 2012; Feinberg et al. 2020).

It is plausible that partisans might then tend to bring similar moral frames to a very wide range of conversations, whether or not a topic has long been explicitly political. For example, polarization in leisure activities and consumer preferences (DellaPosta, Shi, and Macy 2015), though perhaps through different mechanisms. Janoff-Bulman (2009) has argued that although both groups seek to uphold 'community' (perhaps a small step from discussing children), conservatives tend to emphasize *protecting the group* and social order, while liberals emphasize *providing for others* and advancing society through change and the promotion of social justice. Although protection of innocence or defending against physical and/or mental harm may be expressed equally between groups – particularly when talking about children – it may present in starkly different forms. This may occur at least if a concern can be readily linked to topics commonly associated with social order and justice.

This said, parenting style scales, many of which are 'authoritarianism' scales, use parenting because specific questions related to disciplining and nurturing children are thought to tap into views about social hierarchy, autonomy, and conformity (Pérez and Hetherington 2014) or the use of a 'nation-as-family' metaphor in how people think about politics (Lakoff 1996) – not because discussions *about* children are themselves necessarily politically polarized.

Pandemic and Culture War Politics: Through Subsets of Conversations and Over Time in Florida. One possible take-away from past work is that we should expect most people to share similar high-level priorities when it comes to children and yet in the context of *discussing* children, we may still see divergent frames evoked by members of different political parties. Similarly, although there is strong evidence of lifestyle polarization (DellaPosta, Shi, and Macy 2015), such effects seem unlikely to be all encompassing.

That is, although there is likely to be some polarization for at least some topics, we might be less certain how pronounced differing partisan frames might be. Such effects may depend on whether a conversation is more about politics, and associated divisions, than children, and shared experiences and priorities. Further, whether a topic was already polarized prior to the pandemic, and clearly linked to partisan differences in moral values, may influence the extent of polarization on that topic. New topics that may or may not align cleanly on pre-existing divides might be less readily linked to long-standing divisions for most people, even when there is substantial polarization among party activists and highly political interested people (Layman and Carsey 2002). Thus, in considering general versus domain-specific polarization, we therefore evaluate online conversations about education, pandemic policies, partisanship, and culture war issues (relative to general conversations). These topics were particularly salient between 2020-2023 and offer commensurate co-occurrence with **children** or **kids** for comparative study.

We also evaluate the state of Florida as a most likely case for increased salience of pandemic and culture war concerns (relative to the similarly sized by liberal state of New York).

We might expect online conversations in Florida, where Governor Ron DeSantis passed a number high salience laws related to politics and children[3], potentially as part of a planned run for the U.S. presidency based on culture war messaging (DeSantis 2024), to be more polarized than elsewhere. Also, both Florida and New York are distinctive and useful for our analyses because a) their state legislatures and governorships are controlled by a single party (Republicans in Florida and Democrats in New York), making it possible to pass more partisan legislation, b) they are closed primary states, making it possible for us to reliably observe partisan affiliation in voter records, and c) they have large populations, leading to better statistical power[4]. An analysis of

[2]See, for example: https://www.washingtonpost.com/education/2022/10/18/education-laws-culture-war/

[3]https://www.pbs.org/newshour/politics/here-is-a-look-at-the-laws-desantis-has-passed-as-florida-governor-from-abortion-to-guns.

[4]Nebraska and Wyoming have Republican controlled legislatures and governorships as well as closed primaries, but have far

Florida versus New York over time may provide additional information on relationships between any existing polarization of conversations and specific events over the 2019 to 2023 period. This can help adjudicate the extent to which polarization might arise relatively directly from the pandemic and children's welfare concerns versus from the extension of existing partisan conflicts.

Data

For our analyses, we use a previously established, large-scale panel dataset which links Twitter data to publicly available voter records, provided through the vendor TargetSmart (Grinberg et al. 2019; Hughes et al. 2021; Shugars et al. 2021). Approximately 1.5 million Twitter users were matched by unique first and last names to geographical location as of 2018. Due to computational limitations, for the national-level analyses we used a 10% sample of this data set [5]. For the state-level, over-time analyses we retained full samples of Florida and New York rather than a 10% sample.

With this panel, we extracted tweets posted between September 2019 and June 2023 containing terms **children**, **kids**, and **people** (as a comparison). Due to data collection and processing problems in late 2022, we excluded months September through December 2022. These terms are meant to capture direct references of children, and with many fewer misclassifications than terms like **kid** (e.g., "I kid") or **boys** (often referring to sports or athletes). This avoids methodological artifacts in our word embedding analyses, where a tendency of one party to discuss sports more than the other – rather than differences in mentions of children – could drive differences in average word embedding locations. We limited our analyses to users with a provided Democratic or Republican party affiliation (see section Partisanship and gender for more information). Close to 45k of the 150k originally sampled users mentioned **children** or **kids** over the study period, of whom approximately 17k were affiliated with the Democratic or Republican party. These 17k users posted 375k tweets mentioning children or kids. Sample sizes for each analysis, including subset analyses, are included in their respective figures.

For extraordinarily active Twitter users to not influence results far more than others, all word embedding analyses and fightin' words analyses (see Methods section below) consider only one randomly sampled tweet per user. This sampling has the added advantage of reducing computation time, relative to an inversely weighted analysis. It also avoids the need to cluster standard errors by user. Without clustering, we would tend to dramatically overstate the precision of estimates. Analyses with moral dictionaries use user averages rather than tweet sampling. Meaning, we use a user's fraction of tweets with a token matching a moral foundation term.

Partisanship and Gender. To study partisan polarization, we contrasted tweets for Democrats and Republicans. Party affiliations come from voter records for each state, as collected by the voter record vendor. These party affiliations are consistently recorded in states with closed primaries (Hughes et al. 2021), meaning for example that only voters who have registered as a Democrat can vote in a Democratic primary. To simplify our analyses, we exclude other party affiliations and voters with no indicated party preference[6].

We also contrast differences in tweets by gender. Male and female gender has been consistently recorded for all states in the panel, though 1 to 20% of linked user gender across states is unknown / not provided. This male-female gender comparison provides a benchmark for partisan differences in language use.

Text Pre-Processing. We use the quanteda R package (Benoit et al. 2018) for text pre-processing prior to our word embedding analysis. With it, we remove punctuation, symbols, numbers, URL's, hashtags, user mentions, and terms that appeared fewer than 5 times in an analysis subset. We also split hyphens and remove separators (meaning, the unicode separator and control categories). For the moral dictionary analyses, we use the tidytext R package (Silge and Robinson 2016) with its default settings and without special processing to remove hashtags or usernames.

Methods

We measure partisan language use differences in multiple, complementary ways: using moral language dictionaries, fightin' words (Monroe, Colaresi, and Quinn 2008) analysis of moral word use, and word embedding distances. These methods allow us to measure differences on given and previously studied moral dimensions, mine for large differences in language, and measure *any* difference in language use (whether captured by moral dictionaries or not), respectively.

We contextualize the size of partisan language differences using a gender comparison (for word embedding analyses), variation in differences by associated keywords (education, pandemic, and explicitly partisan keywords), and by contrasting trends in Florida and New York over time. Florida and New York are two states of similar size, with closed party primaries which gives us party affiliation data. Each state had single party control of the state legislature and governorship, therefore their conservative (FL) and liberal (NY) governments had distinct rhetoric and policies related to children over this period.

Each of these analyses contributes to our overarching goal of measuring levels and types of polarization in online discussions about children during the 2019 through 2023 period.

Moral Foundation Dictionaries and Party Differences in Moral Language Use. The prevailing source of moral language identification in text is derived from the original Moral Foundations Dictionary (MFD; Graham, Haidt, and Nosek 2009), and it has been widely used (Johnson and Goldwasser 2018; Rezapour, Dinh, and Diesner 2021;

smaller populations than Florida.

[5]Users whose Twitter user ID's ended with 8, leaving roughly 150k Twitter users from the originally matched data – many of whom rarely publicly post tweets.

[6]Either by choice or because a state does not collect the information.

Huang, Wormley, and Cohen 2022; Roy and Goldwasser 2023, e.g.,). Given the paucity of terms in the original MFD (about 320) we use the Moral Foundations Dictionary 2.0 (MFD2; Frimer et al. 2019) for our analyses of moral word use by party. This dictionary provides vice and virtue valence lists of words that are 'highly prototypic' of a given moral foundation (Frimer et al. 2019). Words were chosen through expert selection of candidate words, prototypicality analyses using word embeddings, and crowd-sourced validation (Frimer et al. 2019).

We model differences by party using linear regression, with robust standard errors, on the fraction of a user's tweets which contained any moral term on a dimension. Party affiliation is the independent variable. We use the fraction of tweets containing moral terms to improve the interpretability of estimates. We also chose the MFD2 over the Extended MFD (eMFD; Hopp et al. 2020), which relies on term scores rather than prototypical terms, for this reason. Because the goal of our analysis is inference (and the study of averages across parties) rather than user or tweet-level classification, a larger sample size increases the precision of estimates of the group-level averages and comparisons. These analyses were conducted at the user level rather than tweet level. All users in each analysis were weighted equally, in order to avoid giving increased weight to more active users. We further controlled for the pseudo-log, $ln(x+1)$, of a user's average number of tokens per tweet for each given analysis subset to account for higher probability of dictionary matches due to tweet length alone.

From these models we report the coefficients and 95% confidence intervals for Republican relative to Democrat, with a single moral foundation average as the dependent variable. These estimates are split by vice or virtue terms and by keyword category (see below). We further report sample size, the number of users included in a regression.

Variation in Partisan Language Use by Keyword/Topic. To better understand variation in partisan language use, we ran analyses on subsets of the data, focusing on those more likely to be related to the pandemic or to pre-existing partisan conflicts that appeared to have extended into conversations and political debates about children.

For this, analyses for a given keyword co-occurring in a tweet with **children** or **kids** included four categories: education ("teachers", "students", "schools", "books"), pandemic ("vaccine", "remote", "masks", "distancing"), partisanship/ideology ("republicans", "liberals", "democrats", "conservatives), and what we term partisan flashpoints ("trans", "racism", "migrant", "guns"). We chose terms that would be more unequivocal indicators of a topic, that would not be almost solely used by one party or the other (e.g., "illegals"), and that were relatively frequent in the data. Some more specific phrases, such as "critical race theory" or "CRT", are very rarely used by most Twitter users.

Largest Term-Level Party Differences in Moral Foundations. To evaluate the occurrence of moral terms in child-related tweets most associated with each political group, we used a method for determining the weighted log-odds ratios of terms being chosen by one group over another (Monroe, Colaresi, and Quinn 2008). At a high level, this method identifies terms distinctly used by one group more than another *and* that are relatively frequent. Identified terms are unlikely to randomly occur more times in one group compared to the other. We used the R code (and default prior) provided by its author[7].

Prior to running this, we removed seed terms **children** and **kids** from the analysis. To better understand what moral language is used by each group, we filtered each tweet down to terms that matched the MFD2.

To help interpret these findings, we also ran a bigram version of the same "fightin' words" model. This analysis is intended to disambiguate some term meanings and, in some cases, identify possible term-level misclassification of moral content in this dictionary approach. For it, we included both unigrams and bigrams occurring 10 or more times in the analysis data. We then present the top 2 most polarized bigrams along with their matching most polarized unigrams (if any bigrams for a term occurred 10 or more times).

Polarization in Online Conversations About Children. Current state-of-the-art natural language processing methods have moved beyond dictionary-based methods towards the use of word embeddings (numeric vector representations of words) and other large pretrained language models. More important for this study, recent analyses suggest moral political language may not be adequately captured by dictionaries alone (Kraft and Klemmensen 2024).

To better understand polarization of online conversations about children, we used a modified version of the **conText** embedding regression method described in Rodriguez, Spirling, and Stewart (2021). The method takes all the instances where a specific target term (e.g., **children**) is found in a given corpora, generates GloVe word embeddings from the window of terms surrounding each *context*, and averages across these contexts. For this embedding averaging, we use a window of 6 terms and 200-dimension pre-trained Twitter embeddings [8]. Next, it runs linear regressions across all dimensions of the word embeddings. In our case, party (Republican versus Democrat) is the independent variable and a term's averaged embedding of words in its context window on that dimension as the dependent variable. In its original form, the distance between groups is the Euclidean norm of the regression coefficients. We use a corrected version of the original method due to considerable but easy-to-fix bias, as we explain below. Because coefficients from linear regressions (for a two group comparison) represent differences in averages, this calculates the Euclidean distance between the average embedding locations for each group.

The corrected version of this Euclidean distance estimator removes bias introduced by noise in the estimated regression coefficients (Green et al. 2025). For the squared Euclidean distance, this bias is equal to the sum of the coefficient vari-

[7]https://burtmonroe.github.io/TextAsDataCourse/Tutorials/TADA-FightinWords.nb.html

[8]https://nlp.stanford.edu/projects/glove/, without the *a la carte* embedding discussed in (Rodriguez, Spirling, and Stewart 2021) – and so using associations well before our study period.

ances across dimensions,[9] and so we can simply subtract this value to obtain an unbiased estimate of the squared Euclidean distance (Euclidean distance retains a slight bias) (Green et al. 2025). Our estimates use this corrected value for the *squared* Euclidean norm. Note that although squared Euclidean distance is always positive, these estimates are not. Truncating values at 0 would reintroduce bias, as also noted in other very similar corrections (this bias issue has been regularly noted and similarly corrected across fields; Weir, Wheatcroft, and Price 2012; Nili et al. 2014; Walther et al. 2016). For these estimates, negative values simply indicate little evidence of embedding differences across groups.

Last, as noted by Dodd and Korn (2007) and Green et al. (2025), to our knowledge there is no closed form solution or resampling approach to calculating confidence intervals. Resampling methods provide inaccurate coverage (meaning, a 95% interval calculated from a bootstrap does not actually cover 95% of the sampling distribution). Instead, we provide null distributions from permutation tests, which provide accurate p-values for embedding regression (Green et al. 2025). This follows the originally proposed significance testing for embedding regression (Rodriguez, Spirling, and Stewart 2021), which is unaffected by this correction. It is valid in the case of no regression covariates or clustered data – we sample a single tweet from each user and model using a single independent variable at a time to avoid these complications.

Over-Time Analyses. Last, we consider over-time trends in the embedding regression estimates. This matters because we might expect online conversations about childhood to be tied to specific events, especially the pandemic, or to political debates about policies, the 2020 US presidential election, and 2021 handover of power.

In reporting these analyses, we highlight: the start of the pandemic (in March 2020), the murder of George Floyd and the impetus for ensuing racial justice protests (late May 2020), and the inauguration of Joe Biden as US President (in late January 2021). The last two of these events are major political events not related to the pandemic itself and so more likely to be drive shifts in political language about children by partisan conflict extension than pandemic concerns. We also note emergency use authorizations for use of the COVID-19 vaccine in children, starting (for adolescents) in May 2021. See the Data section for why we focus our analysis on Florida and New York.

Results

Moral Foundations. In Figure 1, we show differences in the fraction of tweets by party that use a term invoking a particular moral foundation. Similar to past research, we find Democrats are somewhat more likely to discuss care and fairness than Republicans. We also see Republicans (though less consistently across the vice and virtue terms on a moral foundation) more likely to invoke sanctity, loyalty, and authority for *some* topics. Republicans invoked fairness in relation to partisan flashpoint keywords more so than Democrats.

Perhaps surprisingly, many of the estimates do not increase in polarity for keywords more strongly associated with the pandemic or with pre-existing political divisions. In particular, *positive* (i.e., virtue) mentions of a moral foundation are relatively consistent across the subsets.

At the same time, we do observe larger differences in specific moral foundations. Republicans were substantially more likely to use *negative* (i.e., vice) terms related to loyalty and authority than Democrats for the pandemic and partisan conversation subsets, and especially so for the partisan subsets. We further see some more ambiguous shifts in other negative terms for other moral foundations, where dimensions that saw higher use by Democrats overall saw more equal or slightly greater Republican use in the pandemic and partisan subsets. Nonetheless, these associations did not strongly differ from mentions of **people**, as we demonstrate in appendix Figure A1 – suggesting only very limited partisan moral framing specific to conversations about children.

Because past work suggests that dictionaries may simply miss meaningful polarization (Kraft and Klemmensen 2024), this result should be considered in tandem with our word embedding analyses below.

Largest Term-Level Differences By Party. Language that differs in moral emphasis often also differs in content. Table 1 reports the results of the fightin' words analyses on moral words – moral terms that tended to be used frequently by one party and far more than by the opposing party. For each term, we list the associated moral foundation, and whether it is categorized as vice or virtue, in the MFD2. And here, we can see differences on race (mentioned more by Democrats), health (Democrats), and religion (Republicans). These appear to reflect well-documented differences in issue priorities and religious affiliation by party.[10] In line with past research to some extent, more of Democrats' terms tended to fall under the care dimension, and Republicans the sanctity and authority dimensions.

Word Embedding Distances. The modified embedding regression provides estimates of squared Euclidean distances in a word embedding space between Democrats and Republicans when talking about children on Twitter. If there is little to no difference in word use between groups, the estimates will be closer to zero (or negative, after bias correction), while a larger estimates indicates a larger difference. As noted in the methods section, the *corrected* distance estimator can produce negative values and truncating these val-

[9] As noted in Green et al. (2025), this method targets an unbiased estimate of the square of a regression coefficient, with population value β^2. However, we only have an *estimate* of β, $\hat{\beta}$, and squaring that gives us the expected value: $E[\hat{\beta}^2] = E[\hat{\beta}]^2 + V[\hat{\beta}]$ (from the definition of variance), where $V[\hat{\beta}]$ is the variance of a regression coefficient and $E[\hat{\beta}]^2$ (in expectation) equal to β^2. Fortunately, standard regression methods provide unbiased estimators for both $\hat{\beta}$ and its variance.

[10] See, for example: https://www.pewresearch.org/politics/2023/06/21/inflation-health-costs-partisan-cooperation-among-the-nations-top-problems/, https://www.pewresearch.org/religious-landscape-study/database/party-affiliation/

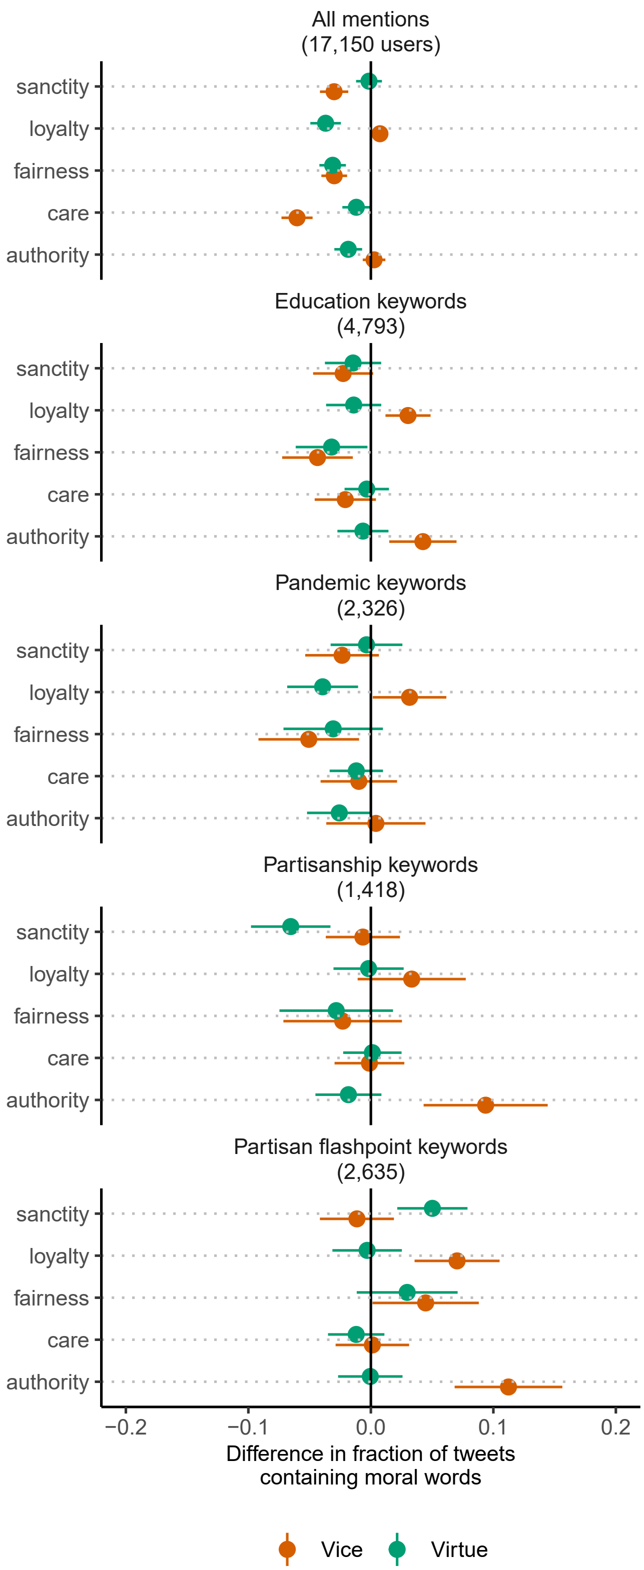

Figure 1: Party moral differences, controlling for differences in tweet length. See text and Figure 2 for the lists of keywords by category.

Democrat	foundation
fuck (the fuck, fuck you)	sanctity (vice)
families (their families, and families)	loyalty (virtue)
care (care if, child care)	care (virtue)
kill (kill children, monsters kill)	care (vice)
health (mental health, the health)	care (virtue)
refuse (refuse to, i refuse)	authority (vice)
racism (& racism, racism thread)	fairness (vice)
horrific	sanctity (vice)
violence (gun violence)	care (vice)
food (food for, food and)	sanctity (virtue)
Republican	**foundation**
jesus	sanctity (virtue)
god (of god, god and)	sanctity (virtue)
blessed	sanctity (virtue)
lord (the lord)	sanctity (virtue)
father (father of, father left)	authority (virtue)
christ	sanctity (virtue)
control (birth control, gun control)	authority (virtue)
order (order to, in order)	authority (virtue)
cheat	fairness (vice)
sexual	sanctity (vice)

Table 1: In **children/kids** tweets: top 10 polarized, moral terms and associated foundations in MFD2. Language that differs in moral emphasis also often differs in content. Beyond moral differences, these terms illustrate greater mentions of race and health, as well as negative emotionality (e.g., expletives, violence, injustice), among Democrats on Twitter, and religious language among Republican Twitter users. Terms in parentheses are the most polarized bigrams for a given term, drawn from a separate fightin' words analysis that included both unigrams and bigrams occurring 10 or more times in the analysis data. No bigrams listed indicates that there were no sufficiently frequent bigrams in the analysis data for a given term.

ues would reintroduce bias in the estimates.

In these models, we compare partisan differences to gender differences, as well as to partisan differences in tweets mentioning **people** (rather than **children** or **kids**). We also analyze overall mentions of children to subsets of mentions that also reference education, pandemic, partisanship, or partisan flashpoints.

Figure 2 displays the results. In this figure, points represent the estimated squared Euclidean distance between a given group and the gray lines represent lower and upper 95% confidence intervals for a null distribution, meaning that points more positive than the gray bars are statistically significant at the 0.05 level. The null distribution is from a permutation test – as noted in the methods section, bootstrapping does not provide accurate confidence intervals for this distance estimator.

We find overall mentions of children are no more polarized by party than by gender. However, education and pandemic conversations are somewhat more polarized, and conversations that reference partisanship or partisan flashpoints are far more polarized.

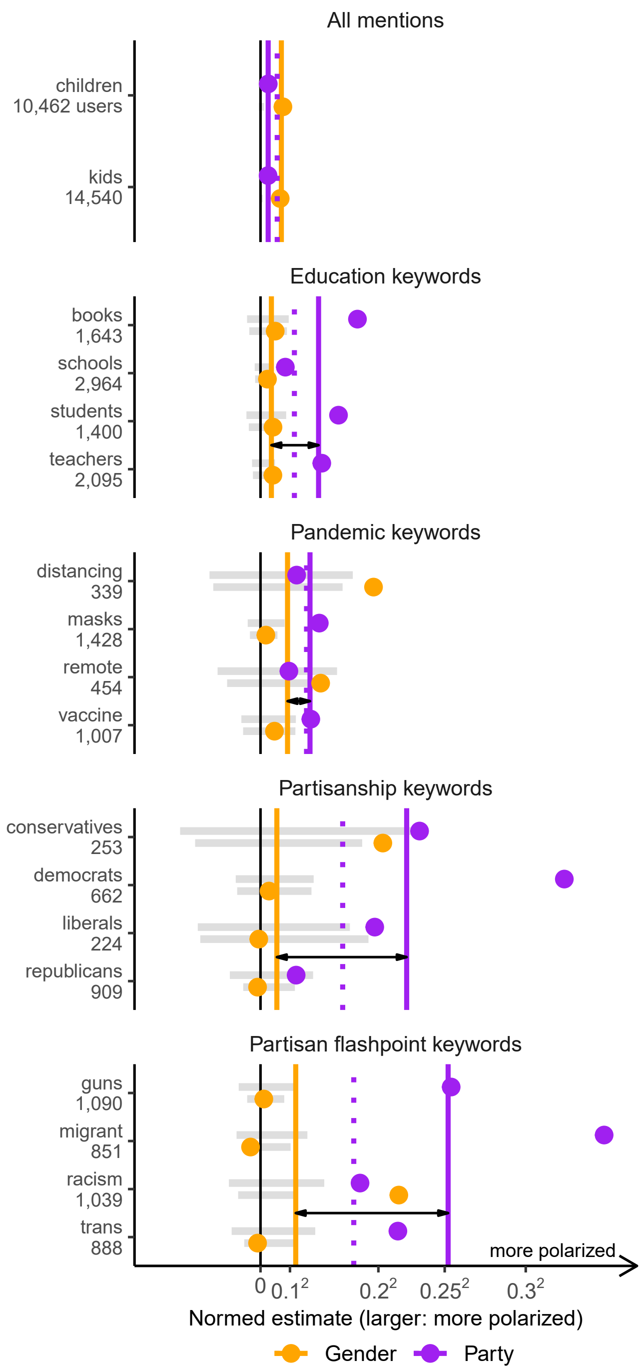

Figure 2: *Embedding Regression by Term.* In this figure, the solid lines indicates the facet (e.g., facet "Education keywords") term-frequency weighted average for terms **children/kids** and the dotted lines facet weighted averages for **people**. Numbers under each term indicate the number of *users* within the 10% sample who used that term at least once, and who are included in the term's embedding regression. Gray bars indicate the 95% confidence intervals for the null distribution of each estimate. These distance estimates can have negative values (see Methods section).

Embedding Distances Over Time. Finally, we examined differences in the use of **children** versus **kids** in a sample of Florida voters. Figure 3 illustrates overall partisan polarization in online conversations about children from late 2019 to mid 2023 in Florida. For the term **children**, we see an increase in polarization in June 2020, just after the murder of George Floyd and several months after the start of the COVID-19 pandemic. Further, we see a large spike in polarization in February 2020 suggesting that conversations about children in Florida may have been polarized before the pandemic as well. For the term **kids**, we observe a moderate increase in polarization in February 2021, just after the presidential inauguration of Joe Biden.

Patterns in New York State are perhaps less clear. Although we seem to observe a similar pattern for the term **kids** in New York as in Florida, we do not observe the same increase in polarization in June 2020 for the term **children**. There is instead an increase in polarization for much of 2021 that, to our knowledge, is not tied to a nationwide political event. Plausibly, it instead aligns with the first emergency authorization for use of a COVID-19 vaccine among adolescents between 12 and 15 in May 2021 [11].

Discussion

This study highlights the limits of moral and political rhetoric in online conversation. We find most Twitter conversations about children among voter file registrants in the United States are not polarized by party, even in the 2020 to 2023 period.

Overview of Results. Conversations related to education and the pandemic were somewhat more polarized, and evoked some divergence in moral frames, while conversations about long-standing partisan divides on race and racism, immigration, gun control, and gender identity were far more polarized. In Florida – where a number of politically contentious laws related to culture war politics and children have been enacted – we also observe somewhat greater polarization following political events than in New York.

Implications. By and large, mentions of children are not polarized by partisanship. That is, when it comes to discussing children in everyday online settings, Democrats and Republicans conversations are mostly similar. In important ways this goes against some recent claims on the increasing politicization of children (e.g., Pulcini et al. 2022) that appear to be overly broad. Conversations mentioning children appear to be divisive when they are about the politics of gender, race, immigration, and guns more than they are about childhood. To the extent that polarized issues become more salient in the everyday lives of children, these pre-polarized issues appear to *stay* just as polarized.

[11] https://www.fda.gov/news-events/press-announcements/coronavirus-covid-19-update-fda-authorizes-pfizer-biontech-covid-19-vaccine-emergency-use (the first EUA for adults was December 11, 2020 https://www.fda.gov/news-events/press-announcements/fda-takes-key-action-fight-against-covid-19-issuing-emergency-use-authorization-first-covid-19)

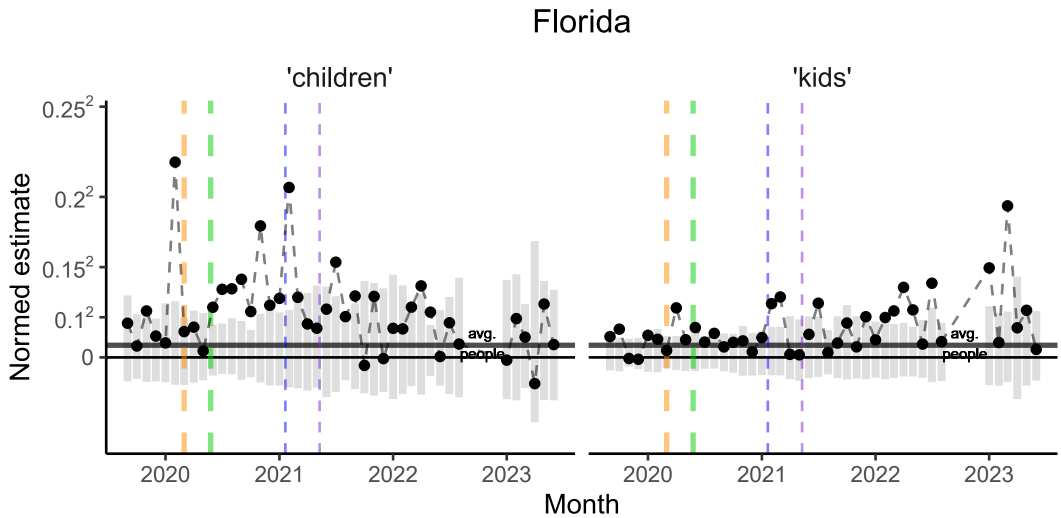

Figure 3: Partisan differences in language use when discussing **children** and **kids**: Florida. Vertical lines indicate the start of the COVID-19 pandemic (orange), the murder of George Floyd (green), the inauguration of Joe Biden as US president (blue), and the FDA's emergency use authorization for the COVID-19 vaccine in children aged 12 to 15 in May 2021 (purple). The horizontal black line indicates the average level of partisan difference in language use when mentioning **people**. Vertical gray bars are 95% confidence intervals for estimates' monthly null distributions from permutation tests. Months September through December 2022 are missing due to data collection problems during that period.

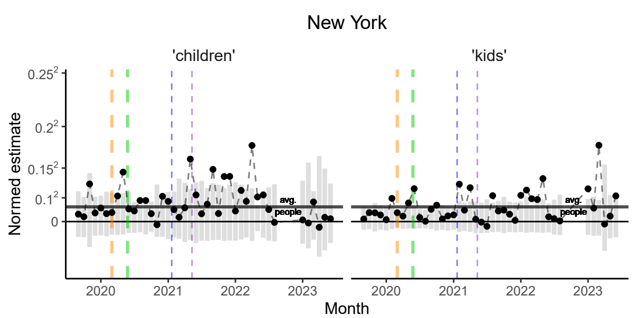

Figure 4: Partisan differences in language use when discussing **children** and **kids** – this figure repeats the analysis in Figure 4 for the state of New York.

These findings have a few implications for thinking about and addressing the effects of this *narrow* polarization. First, for the increasing body of research on emphasizing commonalities for reducing the adverse effects of polarization (Ahler and Sood 2018; Hartman et al. 2022; Syropoulos and Leidner 2023; Levendusky 2023), our findings clearly illustrate that there are far more similarities in how partisans talk about children than not. Democrats and Republicans are not coming from different universes on most children's topics, even with polarized partisan messaging during the COVID-19 pandemic.

Second, however, some issues are political and may stay political, even when they are clearly related to children's welfare (e.g., firearms and firearm safety; McGough et al. 2023). Recognizing that conversations on these topics might remain political allows us to take more seriously scientific research on best strategies for effective communication in polarized settings, and that might not be rooted solely on stating evidence or on appeals to authority. Although there is still a lot that is unknown in this research area and there is no foolproof method for avoiding conflict and misunderstanding, strategies might for example eventually include acknowledging and non-judgmentally listening to experiences and political views (Kalla and Broockman 2020; Ecker et al. 2022; Hartman et al. 2022), rather than merely arguing a point. This could be the case even though many people may prefer these conversations to not include politics at all (Pulcini et al. 2022; Hartman et al. 2022).

Last, although we do perhaps see somewhat greater polarization when 'children' or 'kids' are mentioned in specifically flashpoint and partisan topics rather than 'people', that difference is smaller than other contrasts (e.g., party versus gender). Further research would be needed to assess the robustness of that increased polarization finding, including across longer spans of time and occurring in other political contexts. It is conceivable that this greater polarization reflects *heightened* attempts to misrepresent and vilify political opponents (Wilson, Parker, and Feinberg 2020).

Limitations. Although we establish that morally and politically polarized conversations often did not extend into everyday, public, and online conversations about children, a limitation of this study is that studying polarization in conversations about childhood might tell us little about the polarizing *effects* of political rhetoric and pandemic events on attitudes or on *private* and offline conversations. For example, it is plausible that more polarized conversations, especially ones about some types of moral or out-group content (Brady et al. 2017b; Rathje, Van Bavel, and Van Der Linden 2021), may have been more widely shared and viewed, and so potentially more influential. But even with measures of views, we would still not know whether views of online content translate into attitudes, and there is strong evidence to suggest they often do not (Bail et al. 2020; Guess et al. 2023) (though with some exceptions; see for example Mooijman et al. 2018).

Further, we study only the social media platform Twitter. It is possible that we would observe varied effects in general conversations about children on other platforms. In this, we suspect that we would tend to observe those effects on platforms where most conversation is about politics (e.g., on platforms like Gab with a political and largely far-right user base; Hobbs et al. 2023). Of course, we also do not know how far these findings extend into everyday and *private* interactions. There is now some work to rigorously and ethically conduct such studies (Reeves et al. 2021).

Ethical Statement

Conducting computational social science work allows for us as researchers to examine societal phenomenon at scale, however that does not overshadow the need for thoughtful ethical considerations. In examining such salient topics as morality, politics and concepts of protections or harms to children, it is important that we take concerted care of the data and privacy of users. All efforts to minimize risk or harm and protect the privacy of persons in this study were taken. We do not provide user-related information or associated social media text in order to reduce risk of user identification and to avoid violating Twitter terms of service. See Hughes et al. (2021) for more details on privacy considerations in the original construction of the Twitter panel from public profiles. All data and figures shown in this paper are displayed in aggregate. This work was approved (as exempt human subjects research) by our university's Institutional Review Board (IRB #143475 - University affiliation removed for review). We encourage readers of this work to follow the AAAI ethical guidelines [12] if inspired by the findings.

Acknowledgments

We send many thanks to the Lazer Lab[13] for their contributions to this manuscript.

References

Ahler, D. J.; and Sood, G. 2018. The Parties in Our Heads: Misperceptions about Party Composition and Their Consequences. *The Journal of Politics*, 80(3): 964–981.

Atari, M.; Mehl, M. R.; Graham, J.; Doris, J. M.; Schwarz, N.; Davani, A. M.; Omrani, A.; Kennedy, B.; Gonzalez, E.; Jafarzadeh, N.; et al. 2023. The paucity of morality in everyday talk. *Scientific Reports*, 13(1): 5967.

Bail, C. A.; Guay, B.; Maloney, E.; Combs, A.; Hillygus, D. S.; Merhout, F.; Freelon, D.; and Volfovsky, A. 2020. Assessing the Russian Internet Research Agency's Impact on the Political Attitudes and Behaviors of American Twitter Users in Late 2017. *Proceedings of the National Academy of Sciences*, 117(1): 243–250.

Barker, D. C.; and Tinnick, J. D. 2006. Competing visions of parental roles and ideological constraint. *American Political Science Review*, 100(2): 249–263.

Benoit, K.; Watanabe, K.; Wang, H.; Nulty, P.; Obeng, A.; Müller, S.; and Matsuo, A. 2018. Quanteda: An R Package for the Quantitative Analysis of Textual Data. *Journal of Open Source Software*, 3(30): 774.

Boydstun, A. E. 2013. *Making the News: Politics, the Media & Agenda Setting*. University of Chicago Press.

Brady, W. J.; Wills, J. A.; Jost, J. T.; Tucker, J. A.; and Van Bavel, J. J. 2017a. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28): 7313–7318.

Brady, W. J.; Wills, J. A.; Jost, J. T.; Tucker, J. A.; and Van Bavel, J. J. 2017b. Emotion Shapes the Diffusion of Moralized Content in Social Networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(28): 7313–7318.

Brisbane, L.; Hua, W.; and Jamieson, T. 2023. Morality and the Glass Ceiling: How Elite Rhetoric Reflects Gendered Strategies and Perspectives. *Politics & Gender*, 1–35.

Cahn, N. 2022. The Political Language of Parental Rights: Abortion, Gender-Affirming Care, and Critical Race Theory. *Seton Hall L. Rev.*, 53: 1443.

Deason, G.; and Gonzales, M. H. 2012. Moral politics in the 2008 presidential convention acceptance speeches. *Basic and Applied Social Psychology*, 34(3): 254–268.

DellaPosta, D.; Shi, Y.; and Macy, M. 2015. Why Do Liberals Drink Lattes? *American Journal of Sociology*, 120(5): 1473–1511.

DeSantis, G. G. 2024. 'We believe in education, not indoctrination': Governor Ron DeSantis, critical race theory, and anti-intellectualism in Florida. *Policy Futures in Education*, 14782103241226532.

Dodd, L. E.; and Korn, E. L. 2007. The Bootstrap Variance of the Square of a Sample Mean. *The American Statistician*, 61(2): 127–131.

Donnelly, R.; and Patrinos, H. A. 2021. Learning loss during Covid-19: An early systematic review. *Prospects*, 1–9.

Ecker, U. K. H.; Lewandowsky, S.; Cook, J.; Schmid, P.; Fazio, L. K.; Brashier, N.; Kendeou, P.; Vraga, E. K.; and Amazeen, M. A. 2022. The Psychological Drivers of Misinformation Belief and Its Resistance to Correction. *Nature Reviews Psychology*, 1(1): 13–29.

[12]https://aaai.org/about-aaai/ethics-and-diversity/

[13]https://www.lazerlab.net/

Feinberg, M.; Wehling, E.; Chung, J. M.; Saslow, L. R.; and Melvær Paulin, I. 2020. Measuring moral politics: How strict and nurturant family values explain individual differences in conservatism, liberalism, and the political middle. *Journal of Personality and Social Psychology*, 118(4): 777.

Frimer, J. A.; Boghrati, R.; Haidt, J.; Graham, J.; and Dehgani, M. 2019. Moral Foundations Dictionary for Linguistic Analyses 2.0. *Unpublished manuscript*, available here: https://osf.io/xakyw.

Garten, J.; Boghrati, R.; Hoover, J.; Johnson, K. M.; and Dehghani, M. 2016. Morality between the lines: Detecting moral sentiment in text. *Proceedings of IJCAI 2016 workshop on Computational Modeling of Attitudes*.

Graham, J.; Haidt, J.; and Nosek, B. A. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5): 1029.

Green, B.; Hobbs, W.; Avila, S.; Rodriguez, P. L.; Spirling, A.; and Stewart, B. M. 2025. Measuring Distances in High Dimensional Spaces Why Average Group Vector Comparisons Exhibit Bias, And What to Do about it. *Political Analysis*, 1–8.

Grinberg, N.; Joseph, K.; Friedland, L.; Swire-Thompson, B.; and Lazer, D. 2019. Fake news on Twitter during the 2016 US presidential election. *Science*, 363(6425): 374–378.

Guess, A. M.; Malhotra, N.; Pan, J.; Barberá, P.; Allcott, H.; Brown, T.; Crespo-Tenorio, A.; Dimmery, D.; Freelon, D.; Gentzkow, M.; González-Bailón, S.; Kennedy, E.; Kim, Y. M.; Lazer, D.; Moehler, D.; Nyhan, B.; Rivera, C. V.; Settle, J.; Thomas, D. R.; Thorson, E.; Tromble, R.; Wilkins, A.; Wojcieszak, M.; Xiong, B.; de Jonge, C. K.; Franco, A.; Mason, W.; Stroud, N. J.; and Tucker, J. A. 2023. How Do Social Media Feed Algorithms Affect Attitudes and Behavior in an Election Campaign? *Science*, 381: 398–404.

Hackenburg, K.; Brady, W. J.; and Tsakiris, M. 2023. Mapping moral language on US presidential primary campaigns reveals rhetorical networks of political division and unity. *PNAS nexus*, pgad189.

Haidt, J.; and Graham, J. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social justice research*, 20(1): 98–116.

Haidt, J.; and Joseph, C. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, 133(4): 55–66.

Haidt, J.; Joseph, C.; et al. 2007. The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. *The innate mind*, 3: 367–391.

Hartman, R.; Blakey, W.; Womick, J.; Bail, C.; Finkel, E. J.; Han, H.; Sarrouf, J.; Schroeder, J.; Sheeran, P.; Van Bavel, J. J.; Willer, R.; and Gray, K. 2022. Interventions to Reduce Partisan Animosity. *Nature Human Behaviour*, 6(9): 1194–1205.

Haslam, N.; Rothschild, L.; and Ernst, D. 2000. Essentialist beliefs about social categories. *British Journal of social psychology*, 39(1): 113–127.

Hobbs, W.; Lajevardi, N.; Li, X.; and Lucas, C. 2023. From Anti-Muslim to Anti-Jewish: Target Substitution on Fringe Social Media Platforms and the Persistence of Online and Offline Hate. *Political Behavior*.

Hopp, F. R.; Fisher, J. T.; Cornell, D.; Huskey, R.; and Weber, R. 2020. The extended Moral Foundations Dictionary (eMFD): Development and applications of a crowd-sourced approach to extracting moral intuitions from text. *Behavior research methods*, 53: 232–246.

Huang, X.; Wormley, A.; and Cohen, A. 2022. Learning to Adapt Domain Shifts of Moral Values via Instance Weighting. *Proceedings of the 33rd ACM Conference on Hypertext and Social Media*, 121–131.

Hughes, A. G.; McCabe, S. D.; Hobbs, W. R.; Remy, E.; Shah, S.; and Lazer, D. M. J. 2021. Using Administrative Records and Survey Data to Construct Samples of Tweeters and Tweets. *Public Opinion Quarterly*, 85(S1): 323–346.

Janoff-Bulman, R. 2009. To provide or protect: Motivational bases of political liberalism and conservatism. *Psychological Inquiry*, 20(2-3): 120–128.

Johnson, K.; and Goldwasser, D. 2018. Classification of moral foundations in microblog political discourse. *Proceedings of the 56th annual meeting of the association for computational linguistics (volume 1: long papers)*, 720–730.

Kalla, J. L.; and Broockman, D. E. 2020. Reducing Exclusionary Attitudes through Interpersonal Conversation: Evidence from Three Field Experiments. *American Political Science Review*, 114(2): 410–425.

Kennedy, B.; Atari, M.; Davani, A. M.; Hoover, J.; Omrani, A.; Graham, J.; and Dehghani, M. 2021. Moral concerns are differentially observable in language. *Cognition*, 212: 104696.

Kraft, P. W.; and Klemmensen, R. 2024. Lexical Ambiguity in Political Rhetoric: Why Morality Doesn't Fit in a Bag of Words. *British Journal of Political Science*, 54(1): 201–219.

Lakoff, G. 1996. *Moral Politics: What Conservatives Know That Liberals Don't*. University of Chicago Press.

Layman, G. C.; and Carsey, T. M. 2002. Party Polarization and "Conflict Extension" in the American Electorate. *American Journal of Political Science*, 46(4): 786.

Levendusky, M. 2023. *Our Common Bonds: Using What Americans Share to Help Bridge the Partisan Divide*. University of Chicago Press.

Lin, Y.; Hoover, J.; Portillo-Wightman, G.; Park, C.; Dehghani, M.; and Ji, H. 2018. Acquiring background knowledge to improve moral value prediction. *2018 ieee/acm in-*

ternational conference on advances in social networks analysis and mining (asonam), 552–559.

McGough, M.; Amin, K.; Panchal, N.; and Cox, C. 2023. Child and Teen Firearm Mortality in the US and Peer Countries.

Monroe, B. L.; Colaresi, M. P.; and Quinn, K. M. 2008. Fightin'words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4): 372–403.

Mooijman, M.; Hoover, J.; Lin, Y.; Ji, H.; and Dehghani, M. 2018. Moralization in social networks and the emergence of violence during protests. *Nature human behaviour*, 2(6): 389–396.

Nguyen, T. D.; Lyall, G.; Tran, A.; Shin, M.; Carroll, N. G.; Klein, C.; and Xie, L. 2022. Mapping topics in 100,000 real-life moral dilemmas. *Proceedings of the International AAAI Conference on Web and Social Media*, 16: 699–710.

Nili, H.; Wingfield, C.; Walther, A.; Su, L.; Marslen-Wilson, W.; and Kriegeskorte, N. 2014. A Toolbox for Representational Similarity Analysis. *PLoS Computational Biology*, 10(4): e1003553.

Pérez, E. O.; and Hetherington, M. J. 2014. Authoritarianism in Black and White: Testing the Cross-Racial Validity of the Child Rearing Scale. *Political Analysis*, 22(3): 398–412.

Pew Research Center. 2022a. Parents Differ Sharply by Party Over What Their K-12 Children Should Learn in School. Technical report, Pew Research Center, Washington, D.C.

Pew Research Center. 2022b. Partisans Tend to Cite Different Ideas for What More the Government Should Do For Parents and Children. Technical report, Pew Research Center, Washington, D.C.

Pew Research Center. 2022c. Politics on Twitter: One-Third of Tweets From U.S. Adults Are Political. Technical report, Pew Research Center, Washington, D.C.

Pulcini, C. D.; Berman, R. S.; Raphael, J. L.; and Patel, M. 2022. The Politicization of Children: Implications for Child Health and Public Policy. *Pediatric Research*, 92(2): 345–346.

Rathje, S.; Van Bavel, J. J.; and Van Der Linden, S. 2021. Out-Group Animosity Drives Engagement on Social Media. *Proceedings of the National Academy of Sciences*, 118(26): e2024292118.

Reeves, B.; Ram, N.; Robinson, T. N.; Cummings, J. J.; Giles, C. L.; Pan, J.; Chiatti, A.; Cho, M.; Roehrick, K.; Yang, X.; Gagneja, A.; Brinberg, M.; Muise, D.; Lu, Y.; Luo, M.; Fitzgerald, A.; and Yeykelis, L. 2021. *Screenomics* : A Framework to Capture and Analyze Personal Life Experiences and the Ways That Technology Shapes Them. *Human–Computer Interaction*, 36(2): 150–201.

Reiter-Haas, M.; Kopeinik, S.; and Lex, E. 2021. Studying moral-based differences in the framing of political tweets. *Proceedings of the International AAAI Conference on Web and Social Media*, 15: 1085–1089.

Ren, Z. B.; and Schaumberg, R. 2024. Disagreement Gets Mistaken for Bad Listening. *Psychological Science*, 09567976241239935.

Rezapour, R.; Dinh, L.; and Diesner, J. 2021. Incorporating the measurement of moral foundations theory into analyzing stances on controversial topics. *Proceedings of the 32nd ACM conference on hypertext and social media*, 177–188.

Rodriguez, P. L.; Spirling, A.; and Stewart, B. M. 2021. Embedding Regression: Models for Context-Specific Description and Inference. *American Political Science Review*, 1–20.

Roy, S.; and Goldwasser, D. 2023. "A Tale of Two Movements': Identifying and Comparing Perspectives in #BlackLivesMatter and #BlueLivesMatter Movements-related Tweets using Weakly Supervised Graph-based Structured Prediction. *Findings of the Association for Computational Linguistics: EMNLP 2023*, 10437–10467.

Ryan, T. J. 2014. Reconsidering moral issues in politics. *The Journal of Politics*, 76(2): 380–397.

Shugars, S.; Gitomer, A.; McCabe, S.; Gallagher, R. J.; Joseph, K.; Grinberg, N.; Doroshenko, L.; Foucault Welles, B.; and Lazer, D. 2021. Pandemics, Protests, and Publics: Demographic Activity and Engagement on Twitter in 2020. *Journal of Quantitative Description: Digital Media*, 1.

Silge, J.; and Robinson, D. 2016. Tidytext: Text Mining and Analysis Using Tidy Data Principles in R. *The Journal of Open Source Software*, 1(3): 37.

Silver, J. R. 2020. Binding morality and perceived harm as sources of moral regulation law support among political and religious conservatives. *Law & Society Review*, 54(3): 680–719.

Skar, G. B. U.; Graham, S.; and Huebner, A. 2022. Learning loss during the COVID-19 pandemic and the impact of emergency remote instruction on first grade students' writing: A natural experiment. *Journal of Educational Psychology*, 114(7): 1553.

Syropoulos, S.; and Leidner, B. 2023. Emphasizing Similarities Between Politically Opposed Groups and Their Influence in Perceptions of the Political Opposition: Evidence From Five Experiments. *Personality and Social Psychology Bulletin*, 01461672231192384.

Thornton, C.; Lanyi, K.; Wilkins, G.; Potter, R.; Hunter, E.; Kolehmainen, N.; and Pearson, F. 2023. Scoping the Priorities and Concerns of Parents: Infodemiology Study of Posts on Mumsnet and Reddit. *Journal of Medical Internet Research*, 25: e47849.

Walther, A.; Nili, H.; Ejaz, N.; Alink, A.; Kriegeskorte, N.; and Diedrichsen, J. 2016. Reliability of Dissimilarity Measures for Multi-Voxel Pattern Analysis. *NeuroImage*, 137: 188–200.

Weinzierl, M. A.; and Harabagiu, S. M. 2022. From hesitancy framings to vaccine hesitancy profiles: A journey of stance, ontological commitments and moral foundations. *Proceedings of the International AAAI Conference on Web and Social Media*, 16: 1087–1097.

Weir, J. T.; Wheatcroft, D. J.; and Price, T. D. 2012. The Role of Ecological Constraint in Driving the Evolution of Avian Song Frequency Across a Latitudinal Gradient: Evolution of Birdsong. *Evolution*, 66(9): 2773–2783.

Wilson, A. E.; Parker, V. A.; and Feinberg, M. 2020. Polarization in the Contemporary Political and Media Landscape. *Current Opinion in Behavioral Sciences*, 34: 223–228.

Woodrow, C. 1999. Revisiting images of the child in early childhood education: Reflections and considerations. *Australasian Journal of Early Childhood*, 24(4): 7–12.

Ethics Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? Yes, see background section.
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? Yes
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? Yes, see methods section.
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? Yes, see methods section and conclusion.
   (e) Did you describe the limitations of your work? See – conclusion section.
   (f) Did you discuss any potential negative societal impacts of your work? No.
   (g) Did you discuss any potential misuse of your work? No.
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? Yes, see ethics section.
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? Yes.
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? Yes, see background and methods.
   (b) Have you provided justifications for all theoretical results? Yes, see background.
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? Yes, see background.
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? Yes, see background.
   (e) Did you address potential biases or limitations in your theoretical framework? Yes, see background and methods section.
   (f) Have you related your theoretical results to the existing literature in social science? Yes, see background.
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? Yes, see conclusion.
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? NA
   (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? No – we can release all code but the linked Twitter data cannot be released due to API terms of service restrictions and similar voter record re-release restrictions.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? Yes, and see methods section on error bars.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? No.
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? Yes, see background and conclusion.
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? NA
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   (a) If your work uses existing assets, did you cite the creators? Yes – reference footnotes and citations for the Twitter panel, R packages, and GloVe word embeddings.
   (b) Did you mention the license of the assets? No.
   (c) Did you include any new assets in the supplemental material or as a URL? No.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? No.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? Yes, see ethics section.
   (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? NA

(g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset)? NA

6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...

(a) Did you include the full text of instructions given to participants and screenshots? NA

(b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA

(d) Did you discuss how data is stored, shared, and deidentified? NA

Appendix

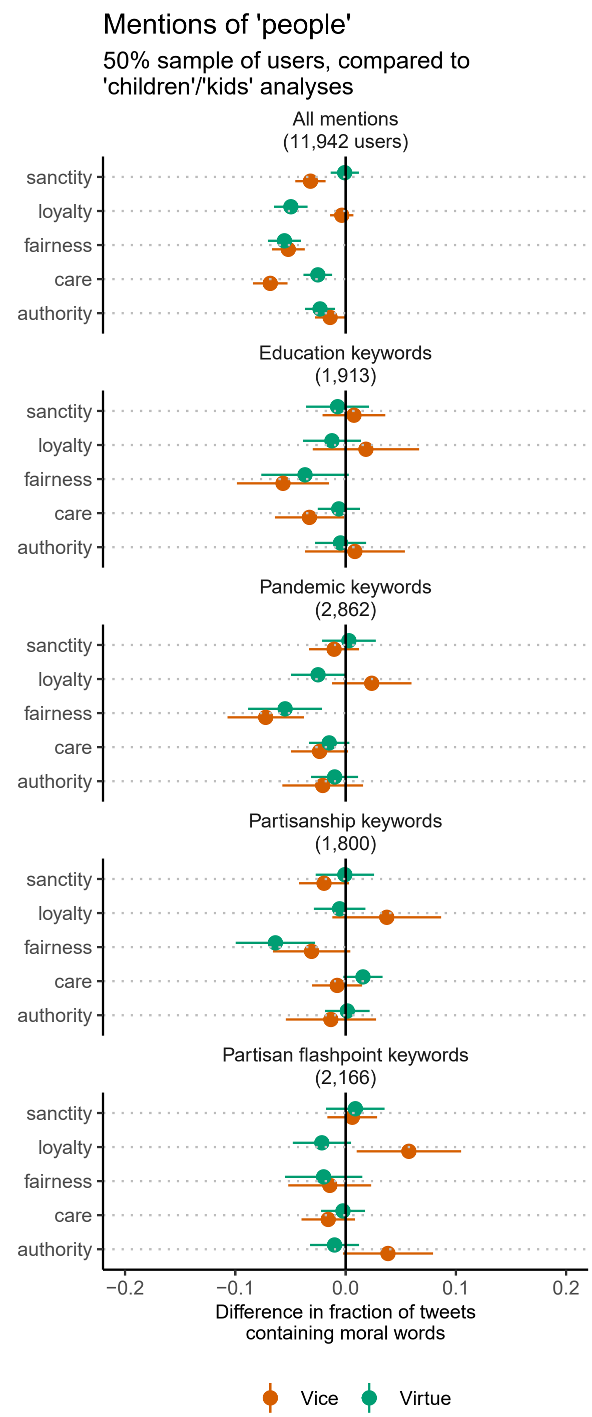

Figure A1: Party moral differences for mentions of **people**, controlling for differences in tweet length. See text and Figure 2 for the lists of keywords by category. Due to computational limitations, we took a further 50% sample of users for this analysis (multiply the sample sizes here by 2 for a frequency comparison with mentions of **children/kids**).