Article:
Heinrich, T, Kobayashi, Y and Long, L (2018) Voters Get What They Want (When They Pay Attention): Human Rights, Policy Benefits, and Foreign Aid. International Studies Quarterly, 62 (1). pp. 195-207. ISSN 0020-8833

https://doi.org/10.1093/isq/sqx081

Voters get what they want (when they pay attention): Human rights, policy spoils, and foreign aid

Tobias Heinrich* Yoshiharu Kobayashi† Leah Long‡

July 27, 2017§

Word count: ≈ 13,000

Abstract

For several decades, scholars have debated whether human rights practices abroad influence states' decisions to distribute foreign aid. This article provides a new solution to this debate by bringing donor government, donor citizens, and recipients' attributes together in a single analytical framework. We argue that donor citizens are more self-serving than previously assumed and do not wholeheartedly support their government punishing human rights abusers that are the sources of important policy spoils. When donor governments anticipate to be accountable to citizens for their policy choices, they make foreign aid decisions to mirror citizens' self-serving policy preferences by avoiding punishing repressive regimes that are the sources of valuable spoils. Our experimental and observational results provide support for our theory. Overall, our findings suggest that aid donors selectively punish human rights violators with aid cuts, but their variegated treatment of human rights violators largely stems from the self-serving policy preferences of donor voters who get what they want.

*Department of Political Science, University of South Carolina, heinrict@mailbox.sc.edu.

†Department of Political Science and International Relations, Nazarbayev University, yoshiharu.kobayashi@nu.edu.kz.

‡Department of Political Science, University of South Carolina, longlj@email.sc.edu.

§Thanks to David Lake, Brandon Gorman, and Charles Seguin for sharing their data sets; to Max Hilbig for help with the survey experiment; and to Beth Bloodgood, Brad Epperly, Carla Martinez Machain, Jim McCormick, participants at the 2016 International Studies Association Annual Convention and 2016 Midwest Political Science Association Annual Meeting, the editors, and anonymous referees for great feedback and comments.

Rhetoric by politicians in liberal democracies suggest human rights to play a prominent role in foreign policy decision-making. In policies such as the allocation of military aid and the use of economic statecraft to military interventions, past research backs up this rhetoric (e.g. Cingranelli & Pasquarello 1985, Efrat 2015, Regan 2000). Given the titular purpose of foreign aid is to improve the lives of people elsewhere, one would expect that human rights practices in recipient countries are also an important concern for distributing such foreign aid. However, empirical research has not been able to settle whether this is the case. Some report that human rights violators receive less aid (e.g. Abrams & Lewis 1993, Cingranelli & Pasquarello 1985), others find no particular relationship (Neumayer 2003*a*,*b*, 2005). Yet, others demonstrate that some human rights violators actually receive more aid (e.g. Carleton & Stohl 1987, Stohl, Carleton & Johnson 1984).

In attempting to resolve the inconclusiveness, recent studies suggest that circumstances affect donors' responses to a recipient's repression. First, the donor government might be unwilling to cut aid to human rights abusers that provide valuable spoils in *quid pro quo* deals for aid. Second, the effect of human rights violations on aid may also depend on media attention to them. The argument goes that news coverage activates voters' sense of morality to which an accountable donor government is incentivized to cater.[1]

These arguments and findings indicate some selectivity by donors and help explain the earlier mixed findings. However, they rest on two different theoretical foundations—one focused on donor-recipient "strategic relationships" and another on the influence of the donor public—, thereby leaving them largely unconnected (at best). We argue and supply evidence that a single theoretical framework can unite these disjoint findings, moving us closer to resolving this long-running debate. Importantly, we also introduce new insights to the literature by deriving and testing implications from our theory which none of the aforementioned perspectives can account for.

[1] For arguments and evidence for the former perspective, see Barratt (2007), Fariss (2010), Nielsen (2013), Esarey & DeMeritt (2016), and Sandlin (2016). The connection with media coverage (and statistical support for it) is presented by Nielsen (2013, 793-794) and Allendoerfer (2015).

We propose a model of foreign aid policy that centers on the accountability relationship between citizens and the government in democratic countries. In this perspective, the locus of policy choices resides in the domestic politics of donors (Bueno de Mesquita & Smith 2012, Milner & Tingley 2015) and with the preferences of the public in particular (Ashworth 2012, Canes-Wrone 2015). Unlike the previous research that assumes an exclusive role of morality in citizens' preferences on foreign aid (Nielsen 2013: 793-794, Dietrich & Murdie 2017: 98, Allendoerfer 2015), we argue that the public also cares about material benefits that aid funds bring about and discounts moral considerations when such benefits are at stake. An original survey experiment fielded in the United States in 2015 corroborates this.

Because (democratic) donor governments are accountable to their constituents, voters' preferences should be reflected in actual aid policy. However, this is not always the case: people's policy preferences matter only when voters are paying attention to these policies and what goes in recipient countries, which is the case when news coverage is high (Van Belle, Rioux & Potter 2004, Eisensee & Strömberg 2007, Baum & Potter 2008, Snyder & Strömberg 2010). Thus, only when coverage by the media is high will the donor government withdraw aid from repressive regimes that provide few policy spoils and refrain from punishing those that are the sources of valuable spoils. Its differential treatment of human rights abusers largely stems from the self-serving policy preferences of voters who get what they want when they pay attention. Furthermore, such trade-offs in the actual aid policy should be less pronounced when news coverage is low as people are not paying attention and thus will not act upon their preferences. Using data on U.S. foreign aid allocations, our statistical tests provide support for these implications.

This article is organized as follows. The first section discusses the debate over the link between human rights and foreign aid, highlighting the two classes of attempts to resolve the debate. While each strand provides evidence for its own hypotheses, we demonstrate that neither can account for the results from the other. By addressing the inconsistency of the earlier literature, this work generates a set of inconsistencies again. The second section

brings the centerpieces of the "strategic" and public opinion approaches together under one coherent framework. We show how our perspective not only accounts for the disjoint findings, but also generates hypotheses linking "strategic factors," news coverage, public opinion, human rights, and foreign aid.

The third and fourth sections study these hypotheses at the micro and macro levels. The third section introduces our survey experiment designed to evaluate the micro-foundations underlying our theory, namely the policy preferences of the donor public. The results support our arguments that people make trade-offs between morality and spoils from aid. The fourth section presents our statistical analysis of the patterns of actual aid allocations. Our main findings are that aid policy reflects the citizens' policy preferences when news coverage is high, but does not do so when attention by the media is low.

Human rights, "strategic interests," news coverage, and foreign aid

One conventional wisdom in the foreign aid literature is that donors' security and economic interests are primary motivations of foreign aid (McKinley & Little 1977, Alesina & Dollar 2000), leaving little room for human rights to matter. In contrast, policymakers in liberal democracies have placed great rhetorical emphasis on the promotion of human rights in other countries. For example, the United States Congress codified such a link in legislation. At first, scholars tried to reconcile this mismatch between words and deeds through (almost exclusively) statistical inquiries. Some studies affirmed that human rights violators are actually punished by receiving less aid from donor states and multilateral agencies (Abrams & Lewis 1993, Cingranelli & Pasquarello 1985, Poe 1992, Apodaca & Stohl 1999, Lebovic & Voeten 2009, Lai 2003). In contrast, other research reported that this is not the case (Neumayer 2003*a*,*b*, 2005) or that (some) human rights violators actually obtained more aid (Carleton & Stohl 1987, Stohl, Carleton & Johnson 1984, Demirel-Pegg & Moskowitz 2009).

This is probably as close as one gets to the ideal-type of "mixed findings."

The recent literature provides several possible reasons for these conflicting findings. Some propose that humanitarian concerns have no place in foreign aid to begin with (e.g. Bueno de Mesquita & Smith 2009). Others argue that research design choices may have produced these varied results.[2] Probably most interesting are two strands of critiques that augment theoretical arguments, considering *when* the donor government would even have incentives to address human right violations.

"Strategic interests." In contrast to earlier work that (implicitly) assumes that aid donors are genuinely interested in responding to human rights violations, recent studies point to a trade-off faced by donors between promoting human rights elsewhere and cutting aid (Barratt 2007, Nielsen 2013, Esarey & DeMeritt 2016, Sandlin 2016). Foreign aid is often provided as a bribe to derive political and economic benefits from recipients, including military cooperation and preferential access to markets (Alesina & Dollar 2000, Bueno de Mesquita & Smith 2009). While addressing domestic ills of recipient countries may bring some positive benefits (by boosting the public image of the donor; see Goldsmith, Horiuchi & Wood 2014), donor governments are willing to set aside human rights considerations when aid cuts would jeopardize benefits from the recipient. This suggests that "strategically important" abusers will not see their aid cut. However, if the recipient provides little benefit, the donor should apply aid cuts to it. Nielsen (2013, 797-799), Esarey & DeMeritt (2016), Farris & Meernik (2015), and Sandlin (2016) report evidence consistent with these arguments.

[2] These are often the sources of debates, such as Carleton & Stohl (1987)'s critique of Cingranelli & Pasquarello (1985) who earlier question previous attempts to estimate the impact of human rights on aid distributions only. The latter point out that it is crucial to analyze the decisions of selecting recipient states (e.g. the "gate-keeping" stage) separately from the disbursement of aid to those states that were selected. In turn, Farris & Meernik (2015, 20-21) and Wang (2016, 483-6) criticize this two-stage modeling choice.

Moral public in donor countries. Others link public opinion and decision-making on foreign aid to explain the lack of systematic evidence for the aid-rights nexus. The argument is that addressing human rights is not obviously intrinsically in the interest of donor governments, except when the public in the donor country demands its government protect human rights of people elsewhere (Murdie & Peksen 2014, Nielsen 2013: 793-794, Peksen, Peterson, & Drury 2014, Allendoerfer 2015). That is, when the public is mobilized, the donor government has incentives to be responsive and to take actions against repressive regimes. As these scholars assume that donor public cares strongly about the moral dimension of foreign policy, the donor government will punish human rights violators by cutting aid. Crucially, however, this should only happen if the public is aware of the human rights violations in the recipient country. Nielsen (2013, 797-9) finds evidence in support of this.

While both approaches can account for the older inconclusive finding, neither can easily explain the results stemming from the other. The "strategic interests" perspective cannot explicate the media-driven impact of human rights because it neglects the role of public opinion. Similarly, it is not straight-forward to reconcile the effect of recipients' "strategic values" with the public opinion perspective as it says little about how benefits from the aid relationships affect aid policies; after all, people are assumed to only care about moral values. The issue here is that these perspectives are based on fundamentally different theoretical underpinnings. The "strategic interest" camp sees the donor as a unitary actor, which Esarey & DeMeritt (2016), Sandlin (2016), and others endow with a richer set of preferences than previous work has done. In contrast, the public opinion-rooted approach assumes two different actors—government and citizens—and suggests that they are in tension with each other. While the government does not care about human rights abroad in this framework, it would react to human rights abuses (presumably begrudgingly) when people pay attention.

Despite the apparent disconnect between these recent perspectives, little has been done to synthesize and integrate their separate arguments and findings, or to evaluate the relative merits of the two. To move the literature forward, research could follow one of (at least)

two paths. First, we could treat these arguments as competing theories and look for additional empirical evidence. Perhaps, one could derive more predictions from each theory and evaluate them. The issue is that evidence supporting one perspective would not necessarily invalidate the other because they have profoundly different theoretical underpinnings. One could sidestep this issue of invalidation by considering the extent to which each theory is applicable (Clarke 2001, Imai & Tingley 2012, Heinrich 2013). Instead, we adopt the second path, which is more theoretically integrative. That is, we synthesize these arguments and findings by considering both "strategic" considerations and the public's influence under a single, new theoretical framework.

In the next section, we present our theoretical argument that centers on political accountability relationships between people and the government.[3] In contrast to the "strategic" perspective, we will arrive at the conclusion that voters fundamentally matter (given a scope condition) and, in divergence from the existing public opinion-centric work, people and the government are not in a contentious relationship with each other. Our theory not only accounts for all the aforementioned statistical results, but also generates novel hypotheses, which a theoretical advancement should do (Bueno de Mesquita 1985, 123).

Accountability and foreign aid policy

The locus of our argument is that policymakers in a (democratic) government are beholden to what the public wants, by and large. That is, policymakers want to be seen as acting consistent with preferences of the public lest they are not reelected. This mechanism sets in when accountability by the public looms large (Snyder & Strömberg 2010, Canes-Wrone 2015). Thus, we first develop an argument for how individuals form preferences about foreign aid decisions and then construct an argument about foreign policy decision-making

[3] Our approach draws on a growing literature on foreign aid that indicates that public opinion about foreign aid is not only consistent and stable (e.g. Noël & Thérien 2008, Paxton & Knack 2011, Milner & Tingley 2013), but also influences policymakers' behavior and actual aid policy (Payaslian 1996, Rioux & Van Belle 2005, Milner 2006, Eisensee & Strömberg 2007, Milner & Tingley 2010).

that is rooted in theories of political accountability. This framework allows us to identify the conditions under which the public matters in foreign policy decisions, connecting research on the relationship between the media, mass preferences, and policy-making (Baum & Potter 2008, Canes-Wrone 2015) specifically and between domestic audiences and foreign policy more broadly (Bueno de Mesquita & Smith 2012).

Individual preferences

We assume the public cares about multiple dimensions of a policy under evaluation. First, the moral implications of aid policy are important for citizens to form policy preferences. This is consistent with many prior arguments on foreign aid, which typically assert that donor citizens view foreign aid (mostly) through a moral lens.[4] As citizens tend to consider financial support to a repressive regime as rendering them complicit in the repression (see Bagnoli 2006), we expect that moral implications of aid policy toward repressive regimes should have a particularly strong impact.

Second, citizens' attitudes toward aid policy may also depend on the evaluation of the security and economic implications as shown in previous work.[5] While foreign aid is often portrayed as a form of charity, it is well known that donor governments use aid to derive economic and security benefits for their citizens (e.g. Alesina & Dollar 2000, Bueno de Mesquita & Smith 2009). In a fundamental way, foreign aid is just like any other policy tool in that it brings tangible benefits to the citizens. Therefore, we expect donor citizens to (also) prefer to give aid to a country that provides economic and security benefits.

The moral and material considerations operate jointly to produce individual preferences over aid policy. However, these two forces do not necessarily work in the same direction

[4] See Milner (2006), Eisensee & Strömberg (2007), Nielsen (2013, 793-794), Heinrich (2013, 425), and Milner & Tingley (2015).

[5] Prior work reports that voters care about material benefits and consequences of foreign policies that are as different as immigration policy (Scheve & Slaughter 2001*a*, Facchini & Mayda 2009), trade policy (Scheve & Slaughter 2001*b*, Mayda & Rodrik 2005), economic sanctions (Heinrich, Kobayashi & Peterson 2017), diplomacy (Tanaka 2016), counterterrorism (Garcia & Geva 2016), and the use of the military (Tomz & Weeks 2013, 856-860).

and often compete with one another. This is the case when donor citizens receive benefits from repressive regimes which were bought via foreign aid. That is, when evaluating aid policy toward repressive regimes, citizens are likely to face a trade-off between moral and material considerations. In fact, existing research suggests that people have to face this trade-off frequently because poor, autocratic countries are particularly conducive to supply policy concessions for the donor in exchange for aid (Bueno de Mesquita & Smith 2009) *and* to abuse human rights (Poe, Tate & Keith 1999).[6]

While the relative saliency of these considerations varies across individuals, we can derive a hypothesis about each dimension on policy preferences when it is conditioned by the realization of the other. In particular, we expect that the effect of moral considerations becomes stronger if there is little economic or security benefit from the aid relationship. If little tangible is at stake, then attention can go to the moral dimension. In contrast, when there is material gain from the aid relationship, an individual may set aside moral qualms and turns a blind eye to aid going to the repressive regime. This brings us to our first hypothesis: *As the benefit in an aid relationship increases, a person becomes less likely to support the withdrawal of foreign aid from that repressive recipient.*

From preferences to policy

Having developed our assumptions about people's preferences on aid policies, we now lay out how and when they translate into actual aid policy. This part of our theory focuses on the electoral connection and draws on the literature on political accountability (Ashworth 2012, Snyder & Strömberg 2010, Canes-Wrone 2015). We assume that the primary goal of the donor government is to survive in office, and this requires support from a large segment of the population in a democratic state (Bueno de Mesquita, Smith, Siverson & Morrow 2003). Therefore, a democratic government has an incentive to do what the public wants.

Although governmental responsiveness ought to be the case for aid policies, it should

[6] This is noted by Heinrich & Kobayashi (2017).

not be perfect. In particular, we would not expect the government to take public opinion into account when it anticipates little public attention and thus little accountability for its policy choices (Eisensee & Strömberg 2007, Snyder & Strömberg 2010, Canes-Wrone 2015). The aforementioned discussion about individual policy preferences suggests that people form opinions based on three pieces of information: human rights practices of the recipient, its security and economic importance to the donor, and its aid relationship with the donor. Therefore, we would expect that aid policy only reflects public opinion when the public is aware of these aspects.

Directly measuring the public's knowledge is difficult and impractical because it would require us to get into individuals' heads and document what they know. Thus, we consider circumstances in which citizens would be informed about all aspects of the aid policy so that they can hold politicians accountable. We argue that this occurs when there is high news coverage about the recipient country *in general*. To see this, consider the supply and consumption of such news in the donor country.[7]

First, when the donor media reports on other (recipient) countries in general, it often provides the audience with information about the economic, diplomatic, and military relationship between the two countries. Given its goal is to attract a larger readership, the news media tries to lay out why the audience should care about the issues. For example, when New York Times wrote about the 2011 presidential election in Uganda, it not only reported about the voting, but also that the incumbent president was "a close American ally

[7] One alternative proxy may be to use news coverage on human rights situations of foreign countries, their importance to the donor, and their aid relationships with the donor. The theoretical motivation for this proxy is that people learn about other countries by reading stories in the newspapers, and thus they tend to be more informed about countries that receive more attention from newspapers. But we believe such news measure is too narrow and based on a rather simplistic view on how people learn and accumulate knowledge. First, the theoretical argument behind this news variable implicitly assumes that people are a passive consumer of news stories and do not seek information themselves. In contrast, people have a variety of means to learn about foreign countries, including talking to peers, watching news on TV, seeking information online, etc. Second, the theoretical model also assumes that all news receive equal attention from people and equally contribute to their knowledge. However, we know that attention to news varies depending on interests in the subject and how the news relate to other topics of interest. For these reasons, we argue that not only news about the aforementioned three aspects matter, but also that news coverage about other aspects of the recipient country matter.

whose relatively small nation gets hundreds and millions of dollars in foreign aid." Similarly, when writing about the Ethiopian Prime Minister Meles Zenawi, the New York Times explained that "[Ethiopia] has a close intelligence and military relationship with the United States government [...] Ethiopia is one of the world's largest recipients of development aid, receiving about $3 billion annually, with the United States one of its largest donors."[8] Along the same lines, the media tends to cover issues that attract viewership, such as scandals, violence, and other information salient to the audience. Since human rights violations are often government-sponsored abuses of civilians, these violent acts catch not only elite media attention (Gorman & Seguin 2015, Table 2), but also see coverage by so-called soft news (Baum 2002, 92-4).

Second, news coverage is not only about the actual written content, but also boosts information-seeking behavior by the audience about the covered topic generally. As people are more exposed to stories about a particular country, they become more attentive to and actively seek information about the country generally (Pelc 2013, 635-6). People are more inclined to make an effort to learn about the country by speaking with peers, searching for more information on the internet, and paying more attention to the country in other news. While not every news story covers our triad of aid, importance, and human rights, greater exposure to stories about the country makes people more interested in, more attentive to, and thus knowledgeable about all of these. Over time, knowledge about a country emerges that may be irreducible to the news that was actually printed. Such attentiveness by people bolsters our argument that larger coverage about the recipient country means that people are more knowledgeable and are thus (more) able to evaluate their own country's policy vis-à-vis the other country.

Taken together, our argument suggests that an election-minded government should follow voters' preferences when the news media informs citizens about the recipient country. That

[8] The sources for the news quotes are Jeffrey Gettleman and Josh Kron, "In Uganda, Unrest Gains Little Ground," *New York Times*, February 18, 2011. Page 4A, and Nicholas Kristof, "What's He Got to Hide?", *New York Times*, January 28, 2012, Page SR11.

is, we expect aid allocation decisions to reflect the aforementioned trade-offs between policy spoils from and morality of foreign aid policy only when people are informed about what goes on in the recipient country and how it stands in relations to one's own country. This leads to the following insight: with a high level of news coverage of the recipient country, we expect that the donor government cuts aid to recipients that are not of high "strategic importance" to the donor country. However, when the recipient is valuable, donors will not cut aid. Our second hypothesis is thus: *when there is a high level of news coverage about the recipient country, the negative effect of human rights violations on foreign aid increases (becomes less negative) as security and economic benefits increase.*

Our theory also gives a prediction for when the scope condition is not met (i.e. when news coverage is too low). When people are unaware, the accountability mechanism will not be at play. Therefore, the trade-offs that the citizens make will not translate into the actual policy decisions. Thus, at a low level of news coverage of the recipient countries, we expect that human rights violations will not affect the aid allocation regardless of the economic or security value of the recipient to the donor. This leads to our third hypothesis: *when there is a low level of news coverage about the recipient country, the effect of human rights violations on aid does not increase as security and economic benefits increase.*

Contrast to previous research

Before turning to our statistical analyses, it is important to emphasize how our theory differs from the existing perspectives. After all, we draw on insights from two previous strands of literature which stand on mostly unrelated theoretical grounds. However, we synthesize their core insights within a single framework that differs from either of the previous views. We demonstrate this by showing how interpretations of the results found in earlier research change under our theory.

First, recent studies by Nielsen (2013), Esarey & DeMeritt (2016), Farris & Meernik (2015), and Sandlin (2016) report that donors refrain from cutting aid to human rights

violators that are important political and economic partners; in contrast, aid to abusive but unimportant regimes gets cut. These scholars see that this evidence is supportive of their "strategic" perspective that focuses on a dilemma faced by the donor government when its aid recipients violate human rights. Such aid cuts would lead to fewer policy perks coming to the donor country. Thus, the selective imposition of aid cuts stems from from a unitary donor (government) with a complex set of preferences. By contrast, our theory such cuts come about via the interplay of two actors, when the donor government has incentives to please the citizens.

Second, Nielsen (2013, 797-799) reports that donors apply aid cuts to recipients when their human rights violations are publicized by the media. Nielsen (2013, 793-794) and Allendoerfer (2015) see this result as support for a perspective emphasizing the influence of a moral public opinion. They argue that the donor public dislikes aid to rights abusers and that people's moral-driven preferences get reflected in aid when the government anticipates the pressure from the public. This latter condition sets in when news coverage is high. While we agree with their general accountability mechanism, our theory suggests that people not only exude morality. Rather, the appreciation of spoils from foreign countries leads citizens to compromise on the morality dimension. Governments, in our view, just obediently cater to these preferences (when people pay attention).

In short, our framework predicts the results from both previous, disjoint strands. In addition, ours offers hypotheses about the interplay of the union of key variables (news coverage, importance, and human rights abuses) from the earlier research. This new insight stems from a premise that people make trade-offs between morality and policy spoils, which we will verify next with a survey experiment.

Analysis of public opinion

We now turn to statistical analyses of individual preferences and the link between those preferences and aid flows. Our first study examines whether donor citizens appreciate tangible returns from an aid-relationship and how these affect their moral considerations. To this end, we design, implement, and analyze a survey experiment.[9]

Experimental setup

We take as given the main result by Allendoerfer (2015) and Heinrich & Kobayashi (2017) that foreign aid to human rights violators draws (very) low support.[10] We recruited 1,414 people via the Amazon's Mechanical Turk (MTurk) platform in September 2015 for a short survey on foreign policy preferences.[11] All respondents were self-reported U.S. residents and

[9] This experimental design is a modification from an earlier, preliminary survey experiment that we carried out about half a year earlier. We provide a full analysis of the earlier survey in Section B in the appendix, and the results are entirely in line with what we present below. That is, the study discussed here is already a replication of the earlier results. Given the recent recognition that published (experimental) research often fails to be replicated (for a prominent discussion, see Open Science Collaboration 2015), we view this as strengthening support for our results.

[10] It is worth mentioning how our experiment differs from that of Allendoerfer (2015). This is particularly important because her findings provide little support for our hypothesis that people sacrifice their moral principles for political and economic benefits from aid relationships. We highlight two key differences in our experimental designs that are possibly consequential for our results. First, even though our subject-matter concerns a sensitive question—support for funding a human rights violator–, which can lead to social desirability bias, and therefore bias the results (Rosenfeld, Imai & Shapiro 2016), Allendoerfer (2015) does not address this issue. Our design does so by using a randomized response design (see below). Second, Allendoerfer (2015, 8-9) uses Chad as a specific recipient country in her vignette whereas we design our experiment more abstractly. The use of actual country names might evoke reactions of the survey-taker that we could not control. In that, our stylization of the scenario in fact poses a "tough test" for measuring opinion formation and changes (see Dickson 2011, 103-109). It is unfortunate that we became aware of her experiment only after we had conducted our survey experiments.

[11] The use of MTurk is controversial among some scholars. We see two advantages to using it. First, it is significantly cheaper to recruit respondents on MTurk so that more surveys can be run in general, which helps with external validity of survey results (McDermott 2011) as well as with replication (see Footnote 9). Second, due to the self-selection of participants and their financial incentives, respondents are expected to pay more attention, increasing the quality of data (Kertzer, Powers, Rathbun & Iyer 2014, 832). That is, treatment non-compliance ought to be rarer.

One potential drawback of using an MTurk-based sample concerns whether our estimates generalize to the target population given the well-known skew of MTurk respondents vis-à-vis the U.S. population (Berinsky, Huber & Lenz 2012, Huff & Tingley 2015, Levay, Freese & Druckman 2016). Some report

were paid \$0.25–0.50 for completing the task (about three minutes) on an author's website. Drawing from the design by Heinrich & Kobayashi (2017), every participant was told that:

> The United States provides foreign aid to help poor countries. In 2016, the U.S. government plans to provide *Amount* to a regime that widely imprisons and tortures members of an ethnic and religious minority.

where *Amount* is randomized between \$25m, \$50m, and \$75m.[12] With equal probability, a respondent sees no additional information (baseline case), a statement that "government officials stress that this foreign aid ensures cooperation with the torturing regime which is vital for U.S. national *Type* interests", where *Type* is either "security" or "economic." These wordings connect the provision of aid to the presence of important benefits from the recipient. They also mesh well with our operationalizations of these concepts in the analyses of aid actual aid flows that we introduce later on.

Asking about the extent of endorsement of foreign aid to human rights violators might make respondents uneasy so that they would not answer truthfully. Similarly, admitting that aid should be used for national interest purposes and not in accordance of morality, nobility, and altruism might engender the same problem. Prior research shows that such social desirability bias is expected to be particularly serious when a survey question involves ethical and moral dilemmas (e.g. Chung & Monroe 2003). We address this source of bias by embedding a randomized response design into our experiment.

---

that estimates from MTurk and nationally representative samples barely differ (Clifford, Jewell & Waggoner 2015), including estimates on foreign policy questions. Further, Hainmueller, Hangartner & Yamamoto (2015) and Wang, Rothschild, Goel & Gelman (2015) demonstrate that adjusting estimates from skewed samples with population data can recover population estimates. Therefore, we report in the appendix the results from a statistical model that includes a full set of demographic variables which we use to post-stratify our results (based on data from the Cooperative Congressional Election Study 2012 (Vavreck & Rivers 2008, Ansolabehere & Rivers 2013)). Our results remain substantively unchanged. In the appendix, we describe in detail this procedure. See Park, Gelman & Bafumi (2004) and Lax & Phillips (2009, 371-372) for post-stratification in general.

[12] The amount of aid is randomized as we wish to control for beliefs over the extent of the U.S.-recipient engagement. If left unspecified, respondents that are told that the recipient is of importance are likely to infer that more aid is given. As cost of aid reduces support (Heinrich & Kobayashi 2017), we would invite (downward) bias from confounding in survey experiments (Dafoe, Zhang & Caughey 2016). We reduce this concern by specifying the aid level.

The basic idea of the randomized response design is to introduce some noise to recorded responses, which in effect make it impossible for anyone else to know individuals' answers to sensitive questions. More specifically, respondents are instructed to use a randomizing device (e.g. coin flips), and depending on the outcome of the randomization, they answer a sensitive question or some unrelated question. Because the researcher has no way of knowing which question each respondent answers, this technique allows the researcher to protect the anonymity of the respondent's answer and thus the respondent would feel free to answer questions truthfully. For example, in one such study, Blair, Imai & Zhou (2015, 1306) instructed respondents to secretly roll a dice before answering a yes-no question. Respondents were told to answer "no" if 1 shows on the dice, to answer "yes" if 6 shows, and to answer a sensitive yes-no question about whether they had social connections to armed groups in civil conflict if 2, 3, 4, or 5 shows. The validation study by Rosenfeld, Imai & Shapiro (2016, 792-6) shows that the randomized response design elicits truthful answers in the presence of sensitive questions.

In our experiment, we ask the respondent to flip a coin and only to evaluate the aid policy (i.e. the randomized vignette from above) if the coin flip shows tails. If the coin lands on heads, the respondent is asked to choose his/her quarter of the year of birth. Under each question, the potential answer–oppose/ support for the aid policy question and quarters of the year for the non-sensitive question–is coded as integers numbering 1–4.[13] The subject checks the appropriate radio button. Because we cannot tell who answered which question, this randomization allows us to protect the anonymity of individual responses. In the aggregate, the proportions of aid-responses (as well their connections to covariates) are identified because we know the population distribution of answers to the unrelated questions

[13] In first case, the question asks, "Do you support or oppose the policy?", with answer options being (1) "strongly oppose," (2) "mildly oppose," (3) "mildly support," and (4) "strongly support." The non-sensitive question reads, "In which month were you born?", with (1) for the first three months of the year, and so forth.

about birth months and the probabilities for the outcomes of coin flips.[14]

As we carry out two distinct analyses in this article, we relegate much detail to Section A in the appendix. There, we provide more information on our survey design, a representative screen shot of the survey instrument, descriptive statistics, the coding of demographic variables, how we approach subjects' shirking (Berinsky, Margolis & Sances 2014), the statistical model to identify the effects, and the Markov Chain Monte Carlo (MCMC) approach to estimate the parameters of the model.

Results

Table 1 provides the posterior summaries for the parameters.[15] The interpretation of estimated parameters is the same as with a conventional ordered probit as our statistical model is designed to remove the noise stemming from the randomized responses (see the appendix for further details). The left-hand column gives the results when we include dummy variables for the two randomized features in the vignette, namely the costs and the interest justifications; the right-hand column model includes the demographic variables. Each number atop is the median of the posterior distribution for the coefficient, and the range in parentheses below gives the 95% central credible interval.

Our central finding is that compared to the baseline of human rights violations in the recipient country without any policy spoils, the support of aid improves if either security or economic interests are invoked. The median estimates for both treatments are positive, and

[14] We provide the following instruction at the beginning to respondents: "[o]n the next page, you will be asked to evaluate a hypothetical situation of the United States providing foreign aid to another country. We understand that opposing or supporting a particular U.S. policy might be sensitive. In order for you to feel free to provide us with truthful answer, we implement an anonymity-protecting mechanism that is commonly used in survey research. We ask you to first flip a coin and then only evaluate the policy if the coin landed on tails. If it landed on heads, proceed to answer an unrelated question. Since only you know which side of the coin showed and what your answer to the unrelated question is, the anonymity is protected. Please take your coin now and flip."

[15] The analysis rests on 1,386 responses after removing shirkers among the survey takers as we explain in detail in the appendix. Shirkers were determined based on information before the treatment was administered.

	Support for foreign aid	
	Bare specification	**With covariates**
Security interests	0.65	0.71
	(0.28, 1.02)	(0.34, 1.09)
Economic interests	0.44	0.45
	(0.09, 0.81)	(0.09, 0.83)
Cost \$25m	0.10	0.07
	(−0.23, 0.44)	(−0.26, 0.41)
Cost \$75m	0.08	0.10
	(−0.26, 0.42)	(−0.26, 0.44)
Age		−0.22
		(−0.38, −0.07)
Gender		0.38
		(0.10, 0.69)
High education		0.02
		(−0.26, 0.30)
Ideology		0.01
		(−0.14, 0.16)
Intercept	−0.6	−0.87
	(−0.97, −0.25)	(−1.35, −0.43)
Cut 1	0.73	0.78
	(0.58, 0.9)	(0.62, 0.96)
Cut 2	1.73	1.85
	(1.29, 2.73)	(1.39, 2.92)
N	1386	1386

Table 1: Posterior summaries for the survey analysis; September data. Each column depicts the results for a separate model. The standalone number is the median posterior estimate, the range below it the 95% central credible interval.

the lower ends of the credible intervals are away from zero. The results are the same when demographic variables are included (right-hand column).[16]

The results support the prediction that security and economic interests mitigate the disapproval of aid going to human rights violators. However, the non-linearity of our model

[16] The median effects of costs are small: compared to aid at \$50m, going either up or down by \$25m changes the support barely. In addition, either effects' credible intervals contain zero. This corroborates findings by Heinrich & Kobayashi (2017) who show that conditional on a human rights violation, costs barely matter.

makes it difficult to get a sense of the magnitudes of those effects and at which level of support these they manifest themselves. Therefore, we calculate predicted probabilities for each support level under the baseline, security interests, and economic interests cases.[17] Figure 1 provides the posterior summaries of the predicted levels of support for aid under each of the three cases. Along the x-axis, we plot the support levels and the y-axis gives the probability that respondents choose each support level. Each dot denotes the median estimate, and the vertical line the 95% central credible interval. The three shades of grey indicate the three different cases: baseline, security interests, and economic interests cases.

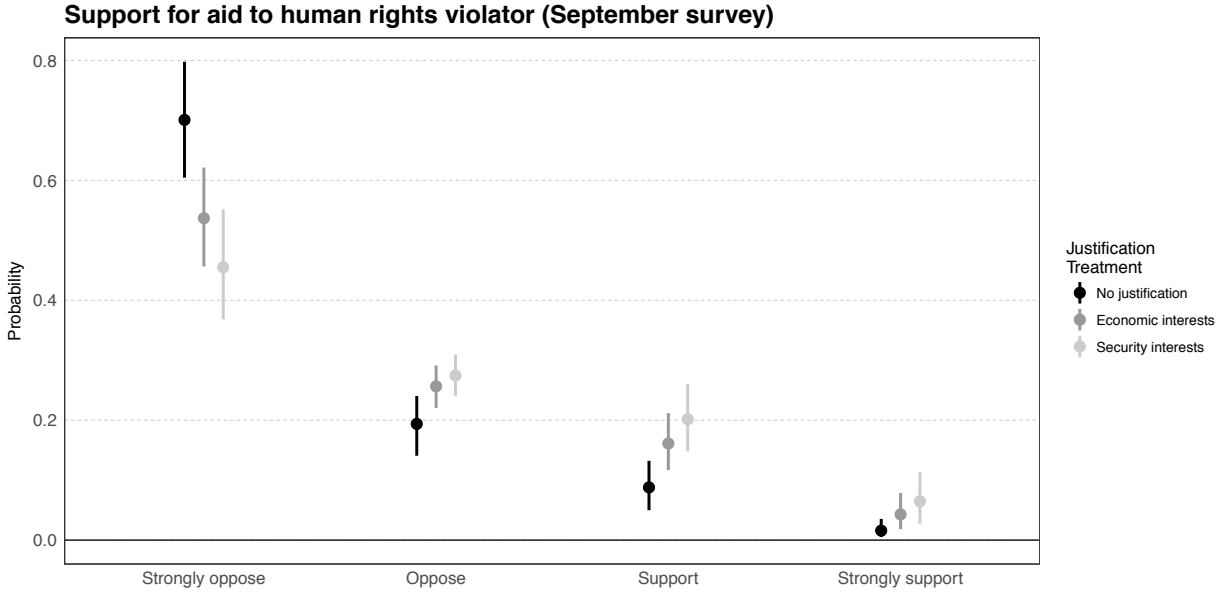

Figure 1: Predicted probability of support levels for aid to human rights violators under different justifications. Each dot denotes the median posterior estimate, and the vertical lines the 95% central credible interval. Black dots and lines denote the probabilities when there was no justification; dark gray counterparts when economic interests are invoked; light gray counterparts when security interests are invoked. Estimates from experimental manipulations only (i.e. left-hand column of Table 1).

Consider first the black dots and lines for the baseline case of human rights violations with neither economic nor security interests invoked. The probabilities that a survey-taker "strongly opposes" or "mildly opposes" the aid to a human rights violator are 0.70 (0.60,

[17] As these quantities depend on all covariates in the model, we marginalize over the cost-related indicator variables for each prediction.

0.80) and 0.19 (0.14, 0.24), respectively.[18] Without further spoils, a recipient's human rights abuses leave a donor's aid allocation very unpopular. However, if we introduce either of the national interests justifications, as indicated by gray dots and lines, strong oppositions wane considerably. When national security justifications are invoked, the strong opposition drops from a median of 0.70 to 0.46 (0.37, 0.55). Furthermore, the median probabilities for "mildly support" and "strongly support" roughly triple compared the baseline; for "mildly support", the probability rises to 0.20 (0.15, 0.26), a noteworthy change up from 0.09 (0.05, 0.13). Results are qualitatively similar when economic justifications are invoked.

In sum, even with spoils, aid to human rights violators does not become wholly embraced by the populace, but its staunch opposition falls considerably and turns moderate. This supports our hypothesis that tangible security and economic benefits can dampen the effect of human rights violations on donor citizens' support for aid.

Analysis of aid flows

With these micro-foundations in hand, we now turn to the second analysis of whether the preferences of donor citizens are reflected in actual foreign aid policies. Our theory suggests that the distribution of foreign aid should follow voters' preferences when these voters are informed and should not when they are not informed.

We base our time-series cross-sectional analysis mostly on the approach and replication data by Nielsen (2013, 795-797) with some modifications. For brevity, a summary of all the variables is in the appendix (Section C).

Empirical strategy

The panel data set by Nielsen (2013, 795-797) includes 21 OECD countries as donors and 113 developing countries as recipients for the period between 1982 and 2004. The unit of

[18] Throughout, we report the median estimate followed by the 95% central credible interval in parentheses.

analysis is the donor-recipient-year. For our purposes, the key variables taken from the dataset by Nielsen (2013) are: the logged total foreign aid per capita for each case (aid commitment in constant U.S. dollars taken from OECD) and the index of human rights violations (Cingranelli, Richard & Clay 2014) (9-point human rights index with 0 indicating full respects for human rights, 8 indicating no respect by the government).

One major change we make to his setup is our operationalization of how much is at stake in a given donor-recipient relationship. To capture donors' so-called "strategic interests," Nielsen (2013, 796) follows the long tradition in the aid allocation literature and uses measures such as the existence of military alliance and the similarity of United Nations General Assembly voting patterns between the donor and recipient (Chapter 3 Neumayer 2005). While these variables are often found to be statistically significant predictors of aid allocations, little theoretical rationale is generally provided for each variable's inclusion; see Farris & Meernik (2015, 3-4) and Esarey & DeMeritt (2016, 9-10) for similar critiques.

Instead, we turn to the work by Lake (2009) on relational contracting between a dominant (donor) and a subordinate (recipient) country. We prefer Lake's measures over a draw from the set of "strategic interest" variables for theoretical as well as empirical reasons.

First, Lake's theory of exchanges in policies between dominant and subordinate states jibes well with our theoretical concept of foreign aid. In our theory, foreign aid is conceptualized primarily as a bribe by the donor government to obtain policy concessions from recipient states. In Lake's theory, the subordinate states give up some of their sovereignty and agree to behave in a way desired by the dominant state in exchange for protection and other assistances such as foreign aid. For example, the United States during the Cold War created and maintained a hierarchy in which many weaker states surrendered some of their autonomy over defense to the United States in return for protection from threats from the Soviet Union and large flows of U.S. financial assistance. Thus, we argue that Lake's theory of dominant and subordinate states dovetails with our concept of aid relationships between the donor and the recipient. Lake's measures capture neatly the benefits that the dominant

(donor) state derives from these relationships as discussed in our theory and invoked in our survey experiment.

Second, Lake's hierarchy measures have been extensively validated within a theoretical context that parallels ours. These measures are indices comprised of economic and security relational variables. While they are in part based on measures that are individually used by other scholars, the resulting indices have been validated to assess whether they capture what they are intended to measure. In fact, Lake (2009, Chapter 3) is devoted entirely to this. Therefore, we can be (more) assured that these indices capture the relational, exchange aspects of foreign aid that we seek to measure.

As great detail about the construction of the security and economic hierarchy indices can be found in his book, we only highlight the contents of them here briefly. His security hierarchy index is the average number of U.S. troops deployed per capita in a subordinate country and the inverse of the number of independent alliances that a subordinate has. The presence of U.S. troops signifies a degree of control and influence, and the more alliances that are not with the United States shows greater independence. The economic hierarchy index is comprised of the inverse of the subordinate state's monetary autonomy and the relative trade dependence with the United States. For example, pegging the exchange rate to the U.S. dollar indicates that U.S. trade and monetary policies will affect, perhaps dominate, the subordinate's economic choices. Similarly, the more relevant trade the subordinate state has with states other than the United States, the less important that state is to the United States. In his thorough validation of the two indices, Lake shows as either hierarchy increases, a subordinate state's policy is more in line with U.S. priorities. This means that as hierarchy increases, the United States is gaining more and more spoils from the relationship; this matches the concepts that we invoked in our survey experiments. We scale both variables so that the minimum and maximum values are at zero and one.

While we gain measures that have a firmer theoretical foundation and have been validated within a context that is close to ours, the cost comes in that the variables only cover recipient

states in relation to the United States. This shrinks our spatial and temporal domain.[19, 20]

Our interest lies in how these hierarchical relations condition the effect of human rights violations on U.S. aid flows. To this end, we specify the predictors by relying on Nielsen's specification, add either the economic or security hierarchy variable, and its respective interaction with recipients' human rights practices.

Further, we need to match our data to our theoretical scope condition, that news coverage of the foreign country is sufficiently high so that people are informed about what goes on in that country and can hold their own government accountable for the chosen policy (Snyder & Strömberg 2010). For our purposes, we would ideally have data on news coverage about politics in each recipient country, but such data are not available. Thus, we use recent data collected by Gorman & Seguin (2015) on how much news coverage in the New York Times specific foreign leaders receive in each year.[21] Data on foreign leaders provide the next-best operationalization. First, this mostly excludes articles about sports and culture which are less informative for our purposes. Second, news articles discussing foreign leaders tend to describe national politics, which include a government's use of repression and its

[19] Because of this, we re-estimate the original main models by Nielsen (2013, 798) to ensure that the basic relationship that he reports (aid decreases as human rights violations increase) holds. We rely on the estimator and approach that we spell out shortly. Table A.1 in the appendix shows the corroborating results.

[20] Readers have pointed out that the gain from using Lake's extensively validated variables comes at too high of a price because of the ensuing sample restrictions; we gain in conceptual precision and credibility, but lose in generalizability. However, even if we had stuck with the unvalidated variables (alliance, UN voting, etc.), we would have had to restrict the sample similarly because our news coverage data, which we will introduce shortly, and Nielsen's news variables are only available for the United States. It is true that Nielsen uses his U.S. news data to proxy news coverage for all donors and thus retains 21 donors in his data, but the news perspective is of lesser importance to his study; it is only the "alternative explanation" (p. 794) in his article. In contrast, the news variable has a fundamental theoretical importance for us. Therefore, the opportunity cost from using Lake's measures in terms of spatial domain is slim given that the important news variable necessitates a similar sample restriction in any case.

[21] The data are counts of New York Times' articles in which a particular leader appears. There can be two leaders in the same country-year, and the same article might cover both leaders. That means that the maximum of the articles for leaders for that country-year would undercount the press coverage, and the sum would overcount it. We proxy the actual news coverage by taking the mean between the maximum and sum, and then apply the logarithm. Results are indistinguishable if we simply use the maximum.

relation to other countries. For example, the newspaper articles mentioned earlier about Ethiopia and Uganda about elections and imprisonment of foreign journalists also mention the country's respective leaders.[22] Finally and most importantly, this broad news variable captures our theoretical argument that recipient countries that receive more coverage about their politics should attract more attention and generate interests from the public, and thus people become more informed about their human rights violations and relationships with the U.S.[23]

To determine the scope condition, we work (for now) only with a subset of the data in which the news coverage that mention foreign leaders is high. For our main analysis, we operationalize this by only using observations with news coverage above the 60^{th} percentile. This captures the cases for which we can be more certain that people are aware of everything that is pertinent to the U.S.–recipient relationship. For the analysis for when the people are uninformed, we mirror the cut-off, relying on cases with news coverage below the 40^{th} percentile.[24]

Like Nielsen (2013), we rely on a Tobit model because our dependent variable is censored at zero. We also make use of recipient-specific random effects, and the parameters are again estimated via MCMC (Hadfield 2010). We rely on locally flat priors; the MCMC shows no signs of non-convergence.

---

[22] Nielsen also introduces a news variable, but it is restricted to articles that mention human rights. While this measure is suitable for these scholars' purposes, our theoretical arguments lead us to focus on knowledge generally (aid, political importance, and human rights). Thus, we use the leader news variable for our main analysis. However, we use Nielsen's data for a robustness check to our data as we discuss later.

[23] Gorman & Seguin's (2015)'s Table 2 shows that leader coverage in the New York Times increases as the respective country's violent conflict increases (which is a well-known crucial predictor of human rights violations by the government), when there is a genocide (which is a drastic but rare manifestation of rights abuses), and as presence of U.S. combat troops (which is part of Lake's measure) increases. This lends support for our claim that the news variable captures the information that people would have to have in order to evaluate U.S. policy.

[24] While our theory tells us that a cut-off is necessary, it cannot tell us where to draw it. Therefore, we probe the robustness by using alternative approaches at the end of the section.

Results

Table 2 provides the results for the key coefficients of our two news subsets and our two hierarchy variables; the results for all coefficients are shown in Table A.2 in the appendix.[25] The left two columns show the results for the high news coverage cases, and the right two columns represent the low news cases; the first and third columns show the economic hierarchy results, and the second and fourth columns give the security hierarchy results. As before, the median estimates and the 95% central credible intervals in parentheses are reported.

	Total aid per capita			
	High news	**High news**	**Low news**	**Low news**
Human rights violations	−0.14	−0.18	−0.11	−0.09
	(−0.29, −0.003)	(−0.31, −0.04)	(−0.29, 0.06)	(−0.23, 0.06)
Economic hierarchy	−1.37		0.70	
	(−3.40, 0.68)		(−1.53, 2.92)	
Violations x economic hierarchy	0.41		0.13	
	(0.001, 0.81)		(−0.55, 0.81)	
Security hierarchy		−2.70		2.50
		(−6.43, 1.13)		(−1.96, 7.00)
Violations x security hierarchy		0.54		−0.65
		(0.02, 1.06)		(−1.70, 0.36)
N	579	728	539	738
Residual SE	3.1	3.05	0.22	1.03
	(1.31, 5.83)	(1.43, 5.41)	(0, 1.12)	(0.15, 2.19)
Random effect SE	3.14	3.25	6.2	5.16
	(2.67, 3.68)	(2.81, 3.75)	(5.26, 7.24)	(4.47, 5.99)

Table 2: Posterior summaries for main models; abbreviated table. Each column depicts the results for a separate model based on the respective hierarchy measure. The standalone number is the median posterior estimate, the range below it the 95% central credible interval. The full results are shown in Table A.2

Consider first the results for the high news coverage subset (left two columns), in which

[25] Section F in the appendix shows formally how we derive all quantities of interest from the estimation equations.

we expect the trade-offs from the public to be present. As the hierarchy variables are strictly non-negative, we can glean the direction of the effects of human rights violations. The negative median coefficients on human rights violations variable suggest that if the hierarchy variables are at their minima of zero, a worsening of the human rights practices leads to less aid. One unit worsening of human rights leads to 14% (0, 29) and 18% (4, 31) cuts in aid when the donor receives few economic and security benefits, respectively.[26] Thus, our evidence reveals that the U.S. punishes rights abusers that are not important.

We test our second hypothesis of whether the marginal effect of human rights violations increases (i.e. becomes less negative) when hierarchy increases. We do this by inspecting the coefficients on the interaction terms by themselves, assuming that the country receives aid.[27] The coefficients on the interaction terms are 0.40 (0.01, 0.80) and 0.53 (0.01, 1.05) for economic and security hierarchy, respectively. At maximum hierarchies (which are scaled to maxima of one), a unit worsening of human rights leads to 40 and 53 percentage-points more aid than when the economic and security hierarchies are at their respective minima. That is, when human rights situations become worse by one unit in a country that is politically and economically important, U.S. aid increases by 26% and 35%, respectively. The results support our hypothesis that there is a trade-off when the news coverage is high: human rights violations lead to less aid when neither economic nor security spoils are at stake, however less so (and perhaps not at all) when the recipient provides policy benefits to the donor.

We also show if, or at which point, these marginal effects of human rights turn from negative to a non-negative. Because this is more easily assessed with graphs, we calculate the marginal effects of the human rights violations for different levels of the hierarchy variables and show them in Figure 2. The x-axis gives the hierarchy measure, the y-axis the percentage change in aid. The solid line provides the median estimate, the dark ribbon spans the 95%

[26] Because the outcome variable is logged, we interpret the coefficients as percentage-changes in aid-per-capita for a unit change in the human rights violation variable.

[27] For further details, see Section F in the appendix.

central credible interval. Consider the left panel for economic hierarchy. When the hierarchy variable is at its minimum of zero, the marginal effect is negative; as hierarchy goes up, the (median) marginal effect increases. The marginal effect crosses zero at around 0.40 and becomes positive on the economic hierarchy, but the lower bound of the credible interval covers zero. A very similar pattern holds for the security hierarchy in the right panel.

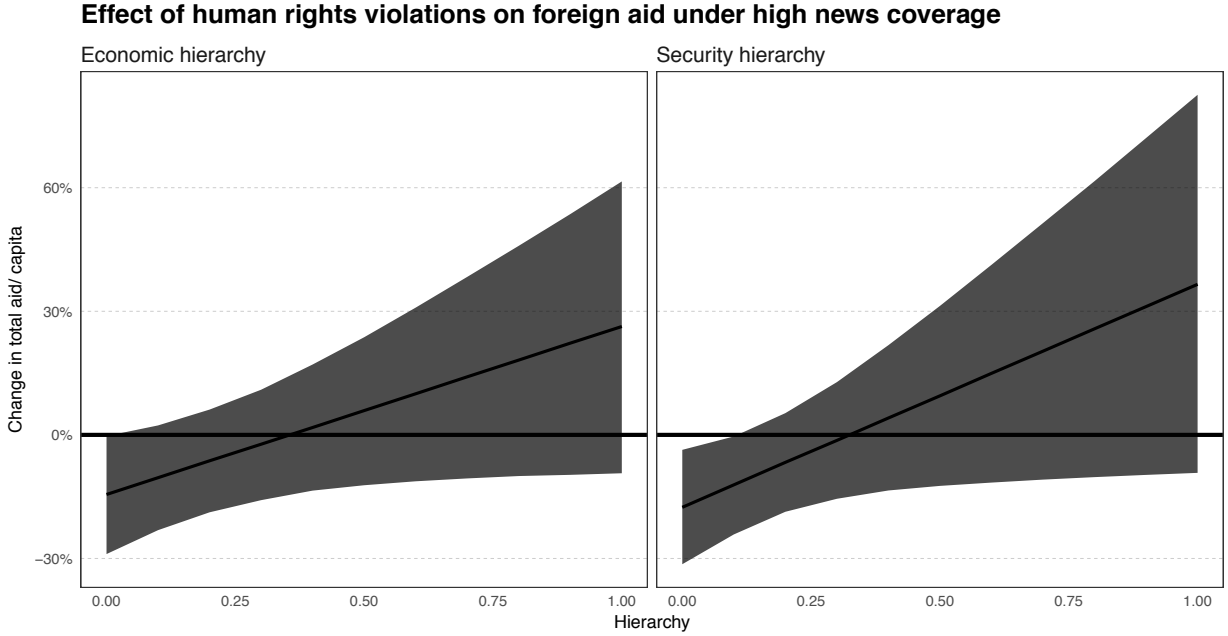

Figure 2: Marginal effect of human rights violations on foreign aid conditional on economic and security hierarchy (high news coverage cases). The x-axis shows the respective levels of the hierarchy variables; the y-axis the percentage change of total aid per capita. The figure is based on models shown in the first two columns in Tables 2 (and A.2). The black line gives the median estimate; the grey ribbon the 95% central credible interval.

Our third prediction is about the relationship between spoils and the effect of human rights violations when people are not knowledgeable about the recipient country. To this end, we turn to the models on the right half of Table 2 which use only data for which the news coverage variable is below its 40^{th} percentile. Assuming that the hierarchy realizations are at their minima of zero, the (median) effects of human rights violations under either model are negative, smaller in absolute value than their high-news counterparts, and their respective credible intervals contain zero. We test whether these marginal effects increase with either hierarchy measure by inspecting the estimated coefficients on the respective

interaction terms. As before, the coefficients on the interaction terms are smaller than those in the model with high news coverage (in fact, negative for security hierarchy), and the credible intervals contain zero. Therefore and as predicted, hierarchy does not systematically increase the effect of human rights violations on aid. These results support our hypothesis that with low news coverage, aid policies do not respond to human rights abuses that reflect people's trade-offs and preferences. Analogous to before, Figure 3 graphs these marginal effects.

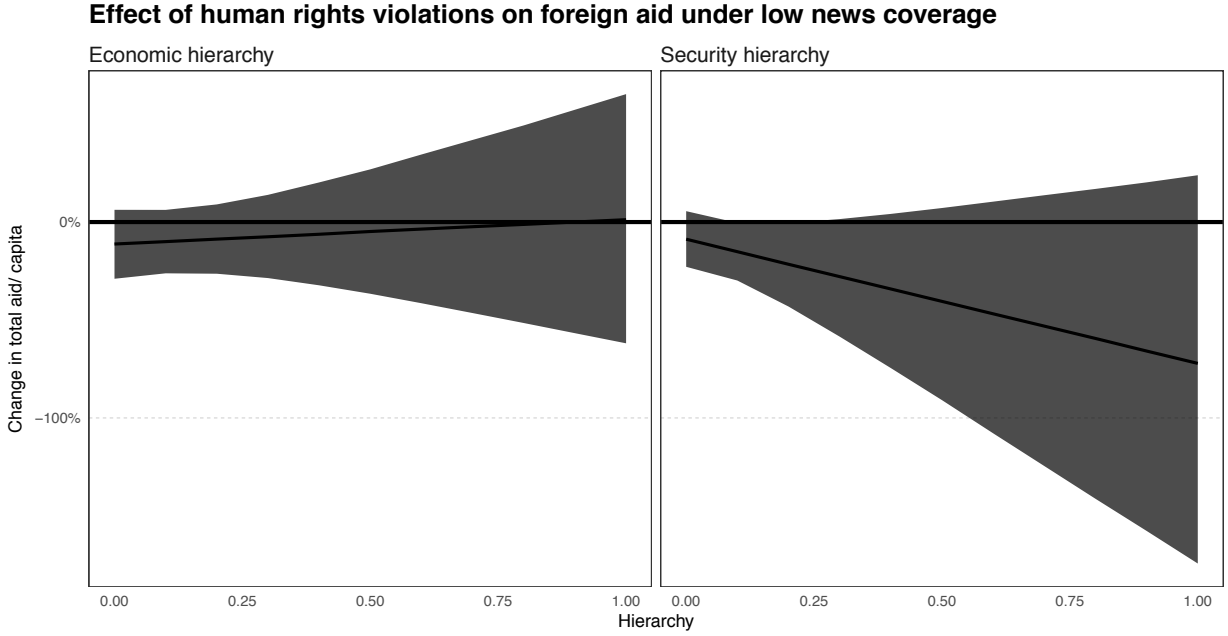

Figure 3: Marginal effect of human rights violations on foreign aid conditional on economic and security hierarchy (low news coverage cases). The figure is based on models shown in the third and fourth columns in Tables 2 (and A.2), but is otherwise constructed analogously to Figure 2.

Finally, we also directly test our hypotheses by calculating the Bayesian posterior probability of support for each. Our first prediction is that for low hierarchy and high news coverage, the marginal effect of human rights abuses is negative. In Figure 4, the left-hand panel depicts the results for marginal effects of human rights abuses when hierarchy variables are at their minima of zero. The black dots denote the posterior probabilities that these statements are true for security (top) and economic (bottom) hierarchy. For both, these probabilities are close to one. Our second prediction is that as hierarchy increases, the effect

of abuses increases (i.e. it becomes less negative). The right hand-side of the figure shows that the two posterior probabilities (black dots) are close to one as well. The gray dots which denote the results for the low-news setting show that the effects are noticeably, sometimes drastically smaller than their high-news counterparts. This provides direct, strong support for our hypotheses.

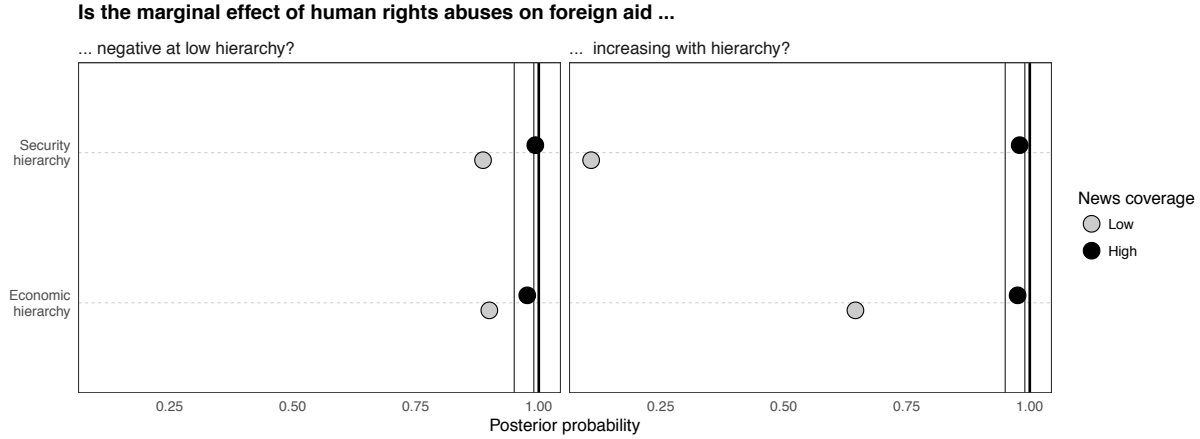

Figure 4: Posterior probabilities in support of hypotheses. The thin vertical lines are at 0.95 and 0.99, respectively; the thicker vertical lines denotes 1.00. The details on how we calculated these numbers are given formally in Section F in the appendix.

Further results

We check the robustness of our results in two ways. First, we employ an alternative source for news coverage. Second, we verify that our subsetting below the 40^{th} and above the 60^{th} percentile of news coverage is not driving the results.

Alternative news source. We argue that the data on country leaders is an effective proxy for news coverage on other countries in general. But it is possible that this data suffers from some systematic bias. In particular, one concern is that our leader data may underreport news about human rights situations of recipient countries while it captures the

security and economic importance of the country.[28] Thus, we augment our analysis by using an alternative measure that more concretely measure news coverage on human rights in foreign countries. We use this as an alternate measure because while it captures human rights violations, it may not capture the more general political relationship of the country to the United States which is an essential concept in our theory. The article by Nielsen (2013, 796) uses the count of news articles in the New York Times that mention human rights in the respective recipient country. We replicate the models presented in Tables 2 (and A.2 in the appendix), but use these more restrictive news data to determine the subsets. When news coverage is high, our main results remain stable (see the upper panel in Figure A.7 in the appendix). When there is little news on human rights practices, hierarchy does not increase the effect of human rights violations, as predicted. In a result that is at odds with our predictions, human rights violations entail a cut in aid at the lowest level of hierarchy even when news coverage is low.

At this point, we can only speculate why this might be the case, but we offer three possible explanations. First, the divergence in the results may stem from the fact that the Nielsen's measure is too restricted and focused too narrowly on information about human rights situations. People can learn about human rights situations through other sources, and thus Nielsen's measure may systematically underrepresent instances in which voters are knowledgeable about human rights violations. This may explain why when using Nielsen's measure, the government still respond to rights violations even in low coverage cases. Second, it is also possible that Nielsen's measure underreports human rights violations in non-U.S. allies but people still learn about them through other sources. The New York Times may

[28] Another concern raised by a referee is that news companies might report less on human rights violations that occur in close allies of the United States. Evidence suggests that the U.S. State Department has tweaked its human rights reports about allies that violate human rights (Qian & Yanagizawa 2009) and that some administrations also actively influenced independent news media, including New York Times, to reduce news coverage on human rights abuses of US close allies (Qian & Yanagizawa-Drott 2017). If this is indeed right, that would mean that high coverage cases using either our leader measure or Nielsen's measure would underrepresent cases of human rights abuses in U.S. allies. However, this would not present a serious challenge to our core empirical findings that in high coverage cases, the donor government disproportionally punish non-allies for their rights abuses.

cover human rights violations less in countries that are not very salient to the United States as it may believe that such news capture less attention from the readers. This can potentially explain the divergent results we see. Finally, it may be that precisely because people are not informed about human rights violations in those cases, other actors such as human rights NGOs step in and pressure the government for an aid cut (for a related argument, see Peterson, Murdie & Asal 2016). We hope that future research addresses this specific scenario.

Alternative scope conditions. As mentioned above, the cutoff below the 40^{th} and above the 60^{th} percentile of news coverage is arbitrary. While using cutoffs better matches our binary theoretical scope conditions and also keeps interpretations of the results manageable, the exact cutoff threshold is not specified by the theory. We deal with this in two ways. First, we simply apply more stringent cutoffs: perhaps above the 80^{th} and below the 20^{th} percentiles of news coverage ensure that people are well informed and not at all, respectively. Second, we can forgo the imposition of a cutoff at the estimation stage and allow the news covariate to modify the effects of rights abuses and of the hierarchy. We include the corresponding triple interaction term (and all constituent terms) and use all the data for the estimation. To make these results comparable to our preferred approach, we subsequently calculate the hierarchy-induced change in marginal effects of human rights abuses on aid by averaging over the analogous news subsets, namely the top and bottom 40%, respectively (for formal statements of this, see Section F in the appendix).

The results using the more stringent cutoffs are depicted in Table A.4 and Figure A.8 in the appendix, and they give actually stronger support than when using the less restrictive subsets. For high-news cases, aid gets cut at the lowest levels of hierarchy, and hierarchy increases these marginal effects. In the low-news setting, aid cuts do not occur at the lowest hierarchy, and more spoils do not increase the marginal effects. Turning to the results computed from the interaction models which we show in Figure A.9, the results

look comparable to our main results from above (with a pronounced steepness as hierarchy increases); the surprising result of aid cuts at a minimum hierarchy and low news coverage remains.[29]

Last, we also show the posterior probabilities in support of the hypotheses using these additional models we just discussed. Figure A.6 in the appendix demonstrates these and it puts in context the anomalies we just discussed. In one of twelve specifications for the high news coverage setting does an estimate drop below 95% posterior probability, and in half of the specifications for the low news setting is the effect of human rights abuses at the minimum hierarchy negative with a posterior probability hovering around 1. This occurs in the models relying on the human rights news variable, the $\leq 40^{th}$ cut-offs, and either hierarchy measure as well as in all interaction models. We hope that future research addresses this, identifying when governments cut aid for unimportant states in the absence of a popular interest for doing so.

Conclusion

For several decades, scholars have struggled to determine whether human rights practices abroad influence donors' decisions to distribute foreign aid. This article provides a new solution to the debate by bringing donor government, donor citizens, and recipients' attributes together in a single framework. Our experimental results suggest that donor citizens not only care about human rights violations abroad but are also benefits-oriented; they hesitate to punish repressive, but valuable recipients. Donor governments, when they are accountable to their citizens, translate such citizens' preferences into aid policy by providing less aid to repressive regimes that are not providing spoils in exchange for aid. However, when there is much benefit at stake, donor governments refrain from conditioning aid on recipients' human

[29] We forgo presenting the coefficient table for the interaction models as triple interactions are notoriously difficult to interpret in them.

rights practices. These complex patterns of aid allocations can explain why earlier studies produced conflicting evidence on the human rights-aid relationship.

At the same time, our theory offers a new way to interpret evidence found in recent research. Several studies report findings that donors remain unresponsive to human rights violations that happen in "strategically important" recipient countries (e.g. Nielsen 2013, Esarey & DeMeritt 2016, Sandlin 2016). They argue that this is because donor *governments* prefer not to jeopardize their beneficial relationships with the recipient. In contrast, our framework and the evidence strongly suggest that it is people that want tangible goods (as well as respect for human rights) and get their government to do what they want when they are attentive.

Our findings have implications for a broader literature on foreign aid. Scholars have increasingly resorted to domestic politics to explain complex patterns of foreign aid decisions, including why donors channel aid through multilateral institutions or NGOs (Milner 2006, Heinrich 2013, Dietrich & Murdie 2017, Milner & Tingley 2015). One premise that many recent theories have in common is that the donor public evaluates foreign aid solely on moral grounds. Our experimental evidence only offers conditional support for this widely-shared assumption. The public evaluates aid policy exclusively through a moral lens only when there is little benefit at stake. This implies that the applicability of the theories is perhaps curtailed; future research will benefit from additional tests of the individual preferences and modifying the theories based on firmer micro-foundations.

Our study also contributes to the broader literature on human rights and foreign policy in two ways. First, our findings complement the existing evidence showing the major role that non-state actors like mass media, human rights organizations, and international organizations may play in bringing the attention to human rights violations (Peksen, Peterson & Drury 2014, Murdie & Peksen 2013, Murdie & Peksen 2013). We also find in the context of foreign aid, news media is important in disseminating information to the public so that people's policy preferences can translate to policy.

Second, our findings help bridge the long-standing gap between the foreign aid literature and the rest of foreign policy studies. Past research gives consistent evidence that human rights play a non-trivial role in a variety of foreign policy decisions, including the decisions to sell arms (Blanton 2000, 2005), allocate military aid (Blanton 1994, Cingranelli & Pasquarello 1985, Poe & Meernik 1995), pursue counterterrorism (Efrat 2015), and militarily intervene (Regan 2000, Choi & James 2016, Murdie & Peksen 2014). However, when it comes to the allocation of foreign aid, prior studies provide a less clear picture on whether human rights are an important concern; a strong skepticism about the relevance of human rights persists. Our findings challenge this pessimistic view: through a fairly nuanced process, human rights considerations profoundly shape foreign aid allocations.