Open Access version of

Dunk, R. J., Goulding, B. W., Fugelsang, J. A., & Friedman, O. (2021). Butt-dialing the devil: Evil agents are expected to disregard intentions behind requests. *Journal of Experimental Social Psychology*, 96, 104188. <https://doi.org/10.1016/j.jesp.2021.104188>

Please note: This version of the paper may differ from the final published article

Butt-dialing the Devil: Evil Agents are Expected to Disregard Intentions Behind Requests

Rebecca J. Dunk¹, Brandon W. Goulding², Jonathan A. Fugelsang², & Ori Friedman²

¹Department of Psychology, York University

²Department of Psychology, University of Waterloo

Author Note

Rebecca J. Dunk (York University); Brandon W. Goulding (<https://orcid.org/0000-0002-4176-4993>); Jonathan A. Fugelsang (University of Waterloo); Ori Friedman (<https://orcid.org/0000-0003-2346-9787>).

Research was supported by a grant from the Social Sciences and Humanities Research Council of Canada awarded to OF.

Please see <https://osf.io/njr28/> for stimuli and data from all experiments. Pre-registrations are available for Experiment 2, <https://aspredicted.org/c9xb4.pdf>; Experiment 3, <https://aspredicted.org/a8u3y.pdf>; Experiment 4, <https://aspredicted.org/dq924.pdf>; and Experiment 5, <https://aspredicted.org/yn3dd.pdf>.

Correspondence concerning this article should be addressed to Ori Friedman, Department of Psychology, University of Waterloo, 200 University Avenue W, Waterloo, Ontario, Canada N2L 3G1. E-mail: friedman@uwaterloo.ca

Abstract

Evil supernatural beings are often depicted as responding to unintended requests, whereas this may be less common in representations of good supernatural beings. This asymmetry suggests that people may expect good and evil agents to differ in their sensitivity to other people's intentions. We investigated this proposal across five experiments on 2231 adult US residents. In Experiments 1 to 4, participants judged whether good or evil agents would grant requests from individuals who varied in their understanding of what they requested, and in whether they executed requests correctly. Across experiments, the good and evil agents were either supernatural beings or regular humans. Participants predicted good agents would be sensitive to intentions behind requests, but predicted evil agents would be comparatively insensitive to these intentions. In some experiments, they also predicted that evil agents would be more sensitive to whether requests were executed correctly. In Experiment 5, participants rated explanations for why an agent would grant a request from someone who did not understand what they were requesting. Participants thought evil agents might grant such requests because they are indifferent to the others' intentions, but participants did not strongly endorse this explanation for good agents. Taken together, our findings suggest that people have distinct expectations of how moral character affects decision-making. They also suggest that people's beliefs about good and evil supernatural beings may be grounded in their views of ordinary humans.

Keywords: moral character; intentions; requests; supernatural beings; prayer

Butt-dialing the Devil: Evil Agents are Expected to Disregard Intentions Behind Requests

In the horror movie *The Gate*, a boy unleashes a horde of demons when he reads a note he does not understand (Takács, 1987). In a medieval folktale, the devil climbs aboard the ark when Noah yells “Enter, enter, you devil of Hell!” to a stubborn donkey (Benkovitz, 1950). Posts on internet forums ask whether you can accidentally sell your soul to the devil (e.g., martrydrebel27, 2013; u/12212888, 2020). These examples illustrate the notion that evil supernatural beings answer summons and requests that were not intended. It is less clear, though, whether anyone thinks it possible to accidentally call on good supernatural beings. People often pray to god and angels, but we hear little about the results of accidental requests to them. Together, this suggests an asymmetry in beliefs about good and evil supernatural beings and their responses to humans.

Here, we ask whether this asymmetry is informative about how people view good and evil agents in general. Recent work suggests people’s moral and social judgments often hinge on assessments of moral character (Goodwin, 2015; Uhlmann et al., 2015). For example, people assign more blame and responsibility for harm to individuals with bad moral character than those with good moral character (Nadler, 2012; Siegel et al., 2017). And while traits like high sociability and competence lead to positive impressions of individuals with good moral character, these traits do not lead to positive impressions of those with bad moral character (Landy et al., 2016). The asymmetry in depictions of supernatural beings may turn out to be informative about the *content* of people’s assessments of moral character.

We suggest the asymmetry may result because people expect moral and immoral agents to differ in their sensitivity to the intentions behind requests. Specifically, people may expect moral agents to consider these intentions, but expect immoral agents to ignore or overlook them.

On this view, evil agents can be accidentally summoned because they care only about the execution of requests, but not whether they are intended. Normally, people do consider intentions when responding to requests. For instance, when asked “please give me the small candle”, people will typically hand over the smallest nearby candle the requester knows about, even if they could instead provide an even smaller one the requester is unaware of (e.g., Epley et al., 2004). Sensitivity to requesters’ intentions is also evident in paternalistic helping. When a requester unintentionally asks for an item that will not help them meet their goal, young children often decline to provide it, and instead hand over another item that can fulfill it (e.g., Martin & Olson, 2013; also see Hepach et al., 2020).

People might expect good supernatural beings to share these cooperative tendencies, and to therefore ignore requests not backed by intentions. People might expect this because they often anthropomorphize supernatural beings, and expect them to think and behave much like humans (e.g., Lane et al., 2016; Lane & Dolins, 2016; Shtulman & Rattner, 2018; see Heiphetz et al., 2016 for a review). For example, although people explicitly claim god can do many things simultaneously, they misinterpret and misremember stories as implying that (like a human), god would have to first finish answering one prayer before responding to another (Barrett & Keil, 1996). In judging what god is like, people are especially likely to attribute human psychological properties (i.e., compared with physical properties; Shtulman & Lindeman, 2016), and they generally expect god to share their beliefs about what is morally acceptable (Heiphetz et al., 2018).

Demons and other evil agents, however, might be viewed as comparatively insensitive to intentions. This view might follow from people conceiving of evil agents as broadly indifferent to norms and other factors that make actions unacceptable (Garrard, 2002; Pettit, 2018), and is

also broadly consistent with findings linking evil with indifference to others' suffering (Gromet et al., 2016). This view of evil also fits with the claim that assessing moral character involves inferring whether an agent's cognitive abilities lead the agent to produce appropriate moral judgments and decisions (Helzer & Critcher, 2018). Intentions are central in most moral judgments (see Young & Tsoi, 2013 for a review), so indifference to them would indicate a major defect in the ability to produce appropriate moral judgments. For example, if an agent were utterly indifferent to intentions, they would be as likely to hinder others as to help them, and would not differentiate between harms inflicted on purpose or by accident.

The Current Experiments

We investigated people's expectations about good and evil agents in five experiments. Some experiments examined judgments about requests directed to supernatural agents, while others asked about requests directed to humans. We asked about both kinds of agents to ascertain whether people's conceptions of good and evil supernatural beings (at least in this context) might be grounded on their views of ordinary humans.

In the first four experiments, participants rated whether agents would grant various requests. Besides manipulating whether these agents were good and evil, we also manipulate two further factors: 1) whether the person making the request understood the request (i.e., and therefore intended it); and 2) how the person executed the request—for example, whether they read words on a note perfectly or with mistakes. An agent sensitive to intentions might be expected to grant imperfectly executed requests. As illustrated by paternalistic helping (Martin & Olson, 2013), even when a request is unclear, the agent might see the goals motivating it (i.e., especially if that agent is viewed as omniscient; Barrett, 2002; Shaman et al., 2016). In contrast, an agent insensitive to intentions might not grant such requests—the agent might not

know what the request was for, or that it was a request at all. Hence, if evil agents disregard the intentions behind requests, their decisions about whether to grant a request might instead depend on surface features, such as how the request was executed.

The final experiment examined people's beliefs about why evil agents might be insensitive to intentions when deciding whether to grant requests. This experiment contrasts three potential explanations: First, evil agents could be viewed as *inattentive* to other's intentions—people might believe they are often too unmotivated or careless to consider others' perspectives. Second, evil agents could be viewed as *incapable* of recognizing others' intentions. Finally, evil agents could be viewed as *indifferent* to intentions. In contrast with the other views, this view holds that evil agents do know what others intend but give this knowledge little weight when responding to requests.

Experiment 1

Methods

Participants. We tested 267 adults ($M_{\text{age}} = 36$ years, range = 22-69, 158 male, 107 female, 2 other). A further 35 participants were excluded for failing at least one of two attention checks or not responding to all test questions. Participants were recruited online using Amazon's Mechanical Turk. They were required to be U.S. residents and to have a HIT approval rate greater than 90%. In this experiment, we recruited 300 participants (150 per between-subjects condition), but we received a few extra responses. All experiments were approved by the Office of Research Ethics at the University of Waterloo.

Materials & Procedure. Participants read about a protagonist who recited a request directed to supernatural beings; see Figure 1 for the script. In one between-subjects condition, the requester asked good beings for a good outcome (angels to protect one's greatest love); in

another between-subjects condition, the requester instead asked evil beings for a bad outcome (demons to harm one's greatest enemy). Instructions above the scenario explained that participants should respond as though they believed the supernatural beings were real (e.g., "Regardless of your beliefs, please respond to the following story assuming that demons are real.").

Participants then rated the likelihood the request would be granted in four circumstances yielded by factorially crossing two within-subjects factors: 1) whether the protagonist understood the request he had made (i.e., and therefore intended it or not) and 2) whether the protagonist recited the request correctly. Participants rated the likelihood the request would be granted using a 5-point Likert scale ranging from *very unlikely* (1) to *very likely* (5). The four items each appeared on a separate row in a matrix grid, along with two items included as attention checks. For example, one attention check read "Marco is tired from the hike. He will go home soon. (Regardless of how you feel, please respond with 'Very Likely')." The four test circumstances appeared in a random order, with the attention checks always second and sixth in the list. The final screen of the survey asked participants to indicate their age in years and their gender. In this experiment, and all others, we have disclosed all measures, manipulations, and exclusions.

Figure 1

Script (left) and Results (right) from Experiment 1.

Marco is hiking through the forest when he notices writing engraved on a large cliff. The writing is a call for **angels to protect one's greatest love/demons to harm one's greatest enemy**. Marco reads the writing out loud.

How likely are **angels to protect Marco's love/demons to harm Marco's enemy** if...

Marco has no idea what the writing means. He recites the words perfectly.

Marco fully understands what the writing means. He recites the words perfectly.

Marco has no idea what the writing means. He recites the words incorrectly.

Marco fully understands what the writing means. He recites the words incorrectly.

Note. Colored text was manipulated between-subjects. Error bars show 95% confidence intervals.

Results and Discussion

We analyzed data from all experiments using generalized estimating equations models (independent correlation matrices) for linear data; analyses of simple effects were Bonferroni-corrected.¹ To our knowledge, sensitivity power analyses have not yet been developed to estimate power for generalized estimating equations, so we are unable to report these here. However, we report sensitivity analysis for analysis of variance (ANOVA) for each experiment in the Supplementary Materials and include the corresponding ANOVA tests there for all experiments, which each return the same pattern of results as those reported here.

The data from all experiments are available online at <https://osf.io/njr28/>. In Experiments 1 to 3, the main analysis included the between-subject predictor *moral character* (whether the

¹ This experiment was not preregistered. We did preregister Experiment 2 to 5, and each preregistration stated the main analyses would use models for ordinal data. We originally analyzed the data that way, but then replaced the analyses with models for linear data. This did not change any of the results.

request granter was evil or good) and the within-subjects predictors *understanding* (whether the requester understood the request) and *accuracy* (whether the request was procedurally correct).

Figure 1 shows participants' mean ratings. All three main effects were significant. Overall, participants judged that requests were more likely to be granted when directed to angels than to demons, Wald $\chi^2 = 9.73$, $p = .002$, when the protagonist understood the request he had made than when he did not, Wald $\chi^2 = 35.26$, $p < .001$, and when he recited it correctly than with mistakes, Wald $\chi^2 = 161.97$, $p < .001$. All three 2-way interactions were also significant: morality*understanding, Wald $\chi^2 = 7.83$, $p = .005$; morality*accuracy, Wald $\chi^2 = 5.56$, $p = .018$; and understanding*accuracy, Wald $\chi^2 = 9.15$, $p = .002$. The 3-way interaction was non-significant, Wald $\chi^2 = 0.36$, $p = .550$.

The interaction between morality and understanding resulted because participants expected angels, but not demons, to be sensitive to whether the protagonist understood the request. Participants viewed angels as more likely to grant the request when the protagonist understood it, $M_{\text{difference}} = 0.59$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.35, 0.83]$, but their expectations for demons did not significantly differ based on the protagonist's level of understanding, $M_{\text{difference}} = 0.21$, $SE = 0.10$, $p = .198$, $CI_{95\%}[-0.05, 0.47]$. Also, when the protagonist understood, angels were judged more likely than demons to grant requests, $M_{\text{difference}} = 0.52$, $SE = 0.12$, $p < .001$, $CI_{95\%}[0.21, 0.82]$, but when he did not understand the writing, ratings for angels and demons did not significantly differ, $M_{\text{difference}} = 0.14$, $SE = 0.13$, $p > .999$, $CI_{95\%}[-0.21, 0.49]$.

The interaction between morality and accuracy resulted because participants viewed accuracy as more important for demons than angels. Participants expected both kinds of beings would be more likely to grant correctly recited requests, but this effect was larger for demons, $M_{\text{difference}} = 1.31$, $SE = 0.14$, $p < .001$, $CI_{95\%}[0.95, 1.67]$, than for angels, $M_{\text{difference}} = 0.90$, $SE =$

0.11, $p < .001$, $CI_{95\%}[0.62, 1.18]$. In addition, angels and demons did not significantly differ in their likelihood of granting correctly recited requests, $M_{\text{difference}} = 0.12$, $SE = 0.14$, $p > .999$, $CI_{95\%}[-0.24, 0.49]$, but angels were more likely than demons to grant incorrectly recited requests, $M_{\text{difference}} = 0.53$, $SE = 0.13$, $p < .001$, $CI_{95\%}[0.18, 0.88]$.

Finally, the interaction between understanding and accuracy resulted because although requests were viewed as more likely to be granted if recited correctly than incorrectly at both levels of understanding, the advantage for correct reciting was greater when the protagonist understood the request, $M_{\text{difference}} = 1.23$, $SE = 0.10$, $p < .001$, $CI_{95\%}[0.97, 1.49]$, than when he did not, $M_{\text{difference}} = 0.98$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.73, 1.23]$.

These findings suggest that people expect moral and immoral agents to differ in their sensitivity to the intentions behind requests. Angels were judged more likely to grant the protagonist's request when he understood the request. In contrast, expectations for demons did not significantly differ based on whether the protagonist understood or not. We found the opposite pattern, though, when examining how judgments were affected by how the protagonist executed the request—demons were expected to care more than angels about whether the protagonist recited the request correctly or with errors.

In manipulating moral character, however, we varied not only the nature of the agent granting the request but also the content of the request itself. The request directed to angels was for a good outcome, whereas the request directed to demons was for a bad one. The next experiment tests whether the findings replicate if the content of the request is not specified. Further, to examine the robustness of the effect the next experiment uses a cover story where requests are directed to god or the devil, rather than to angels and demons.

Experiment 2

Methods

This experiment was preregistered at <https://aspredicted.org/c9xb4.pdf>.

Participants. We tested 336 adults ($M_{\text{age}} = 37$ years, range = 20-69, 211 male, 124 female, 1 other). A further 67 participants were excluded for failing at least one of two attention checks or not responding to all test questions. Participants were recruited online using Mechanical Turk. They were required to be U.S. residents and to have a HIT approval rate greater than 90%. In this experiment and the next, we recruited 400 participants (200 per between-subjects condition), but again received a few extra responses. We increased the target from the first experiment to reduce the chance of failing to detect a 3-way interaction (should one exist).

Materials & Procedure. This experiment had the same design as the first experiment but used a different cover story where a person reads aloud a note found in a library book (see Figure 2). The note directed a request to either god or the devil (manipulated between-subjects) and instructions above the story asked participants to respond as though they believed in the corresponding being (e.g., “Regardless of your beliefs, please respond to the following story assuming that God is real.”) In contrast with the first experiment, the story did not specify content of the request, or whether it was for a good or bad outcome. After reading the story, participants again rated the likelihood the request would be granted in four circumstances varying in whether the protagonist understood what he had requested, and whether he recited the request correctly (both factors manipulated within-subjects).

Figure 2

Script (left) and Results (right) from Experiment 2.

Jordan finds a note in a library book. The writing on the note is a request to God/the Devil. Some of the words are strange and difficult to read. Jordan decides to read the note out loud.

How likely is God/the Devil to grant the request if...

Jordan has no idea what the writing means. He recites the words perfectly.

Jordan fully understands what the writing means. He recites the words perfectly.

Jordan has no idea what the writing means. He recites the words incorrectly.

Jordan fully understands what the writing means. He recites the words incorrectly.

Note. Colored text was manipulated between-subjects. Error bars show 95% confidence intervals.

Results and Discussion

Figure 2 shows participants' mean ratings. All three main effects were again significant.

As before, participants judged that requests were overall more likely to be granted when the protagonist understood what he had requested than when he did not, Wald $\chi^2 = 18.41$, $p < .001$, and when requests were recited correctly than with mistakes, Wald $\chi^2 = 91.73$, $p < .001$.

However, the direction of the main effect of moral character was opposite to the pattern in Experiment 1. Now judgments that requests would be granted were overall higher for requests directed to the devil than to god, Wald $\chi^2 = 5.44$, $p = .020$. We also replicated all interactions from the first experiment. As before, all 2-way interactions were significant:

morality*understanding, Wald $\chi^2 = 5.38$, $p = .020$; morality*accuracy, Wald $\chi^2 = 14.37$, $p < .001$; and understanding*accuracy, Wald $\chi^2 = 10.44$, $p = .001$. The 3-way interaction was again non-significant, Wald $\chi^2 = 0.03$, $p = .871$.

The interaction between morality and understanding resulted because participants expected god, but not the devil, to be sensitive to the protagonist's understanding. They thought god would be more likely to grant the request when the protagonist understood it than when he did not, $M_{\text{difference}} = 0.45$, $SE = 0.10$, $p < .001$, $CI_{95\%}[0.19, 0.7]$, but expectations for the devil did not significantly differ based on the protagonist's understanding, $M_{\text{difference}} = 0.13$, $SE = 0.09$, $p = .966$, $CI_{95\%}[-0.12, 0.38]$. Also, when the protagonist understood, participants gave similar ratings for god and the devil, $M_{\text{difference}} = -0.08$, $SE = 0.13$, $p > .999$, $CI_{95\%}[-0.42, 0.25]$, but when the protagonist lacked understanding, they rated the devil more likely than god to grant requests, $M_{\text{difference}} = 0.40$, $SE = 0.12$, $p = .005$, $CI_{95\%}[0.08, 0.71]$

The interaction between morality and accuracy resulted because participants viewed accuracy in reciting requests as more important for the devil than god. Both agents were viewed as more likely to grant correctly than incorrectly recited requests, but this effect was larger for the devil, $M_{\text{difference}} = 1.00$, $SE = 0.12$, $p < .001$, $CI_{95\%}[0.68, 1.33]$, than god, $M_{\text{difference}} = 0.43$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.21, 0.66]$. Also, the devil was viewed as more likely than god to grant correctly recited requests, $M_{\text{difference}} = 0.52$, $SE = 0.13$, $p < .001$, $CI_{95\%}[0.18, 0.87]$, but they did not significantly differ in judgments about granting incorrectly recited requests, $M_{\text{difference}} = 0.05$, $SE = 0.12$, $p > .999$, $CI_{95\%}[-0.28, 0.37]$.

Finally, the interaction between understanding and accuracy resulted because although participants judged that requests were more likely to be granted if recited correctly than incorrectly at both levels of understanding, the advantage for correct reciting was greater when the protagonist understood, $M_{\text{difference}} = 0.85$, $SE = 0.08$, $p < .001$, $CI_{95\%}[0.65, 1.06]$, than when he did not, $M_{\text{difference}} = 0.58$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.34, 0.83]$.

In sum, participants' judgments again depended on whether requests were directed to good or evil beings; this factor was more tightly controlled than in the first experiment as the vignette did not specify (or vary) the content of the request. As in the first experiment, participants expected the good agent would privilege requests when the protagonist understood what he was requesting, whereas they did not expect the evil agent to care about this. Also replicating the first experiment, participants judged that the evil agent would care more than the good one about whether requests were recited correctly.

The findings so far suggest people have differing expectation about good and evil supernatural beings. But the findings do not tell us whether these differing expectations apply to good and evil agents in general. To explore this question, we next investigated judgments about requests addressed to ordinary humans.

Experiment 3

Methods

This experiment was pre-registered at <https://aspredicted.org/a8u3y.pdf>.

Participants. We tested 319 participants ($M_{\text{age}} = 37$ years, range = 20-69, 198 male, 121 female). A further 87 participants were excluded for failing at least one of two attention checks or not responding to all test questions. Participants were recruited online using Mechanical Turk. They were required to be U.S. residents and to have a HIT approval rate greater than 90%.

Materials & Procedure. This experiment has the same design as the previous experiments but used a new scenario about a bureaucrat whose job is to decide whether to approve requests from the public (see Figure 3). The bureaucrat was either good or evil (manipulated between-subjects) and participants rated the likelihood of the bureaucrat granting four requests varying in whether they were from people who understood what they requested,

and in whether the corresponding paperwork had been filled out correctly (both factors manipulated within-subjects).

Figure 3

Script (left) and Results (right) from Experiment 3

Pat is a bureaucrat who reads requests from the public. His job is to decide whether requests will be approved or denied. Pat is exceptionally **good and moral**/**evil and immoral**. This greatly affects his decisions about whether to grant or deny requests.

How likely is Pat to grant a request from...

A person who is obviously confused about what they are requesting. The person filled out the paperwork correctly and made no mistakes.

A person who is obviously confused about what they are requesting. The person filled out the paperwork incorrectly and made many mistakes.

A person who obviously understands what they are requesting. The person filled out the paperwork incorrectly and made many mistakes.

A person who obviously understands what they are requesting. The person filled out the paperwork correctly and made no mistakes.

Note. Colored text was manipulated between-subjects. Error bars show 95% confidence intervals.

Results and Discussion

Figure 3 shows participants' mean ratings. The three main effects were again significant. As in the first experiment, requests were judged as more likely to be granted when directed to someone good rather than evil, $\text{Wald } \chi^2 = 42.40, p < .001$. And as in both previous experiments, requests were judged more likely to be granted when protagonists understood what they were requesting than when they did not, $\text{Wald } \chi^2 = 70.58, p < .001$, and when requests were made correctly than with mistakes, $\text{Wald } \chi^2 = 109.55, p < .001$. As in both previous experiments, all three 2-way interactions were significant: morality*understanding, $\text{Wald } \chi^2 = 44.73, p < .001$;

morality*accuracy, Wald $\chi^2 = 9.34$, $p = .002$; and understanding*accuracy, Wald $\chi^2 = 44.15$, $p < .001$. The 3-way interaction was again non-significant, Wald $\chi^2 = 0.06$, $p = .815$.

Consistent with Experiments 1 and 2, the interaction between morality and understanding resulted because participants thought the good bureaucrat, but not the evil one, would be sensitive to the protagonist's understanding. Participants expected the good bureaucrat would be more likely to grant requests if protagonists understood what they were requesting, $M_{\text{difference}} = 0.95$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.73, 1.18]$, whereas expectations that the evil bureaucrat would grant requests did not significantly differ based on the protagonist's level of understanding, $M_{\text{difference}} = 0.11$, $SE = 0.09$, $p > .999$, $CI_{95\%}[-0.14, 0.35]$. Also, when the protagonist understood, the good bureaucrat was viewed as more likely than the evil one to grant requests, $M_{\text{difference}} = 1.01$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.77, 1.25]$, but when the protagonist did not understand, ratings for good and evil bureaucrats did not significantly differ, $M_{\text{difference}} = 0.16$, $SE = 0.13$, $p > .999$, $CI_{95\%}[-0.17, 0.49]$.

Inconsistent with Experiments 1 and 2, the interaction between morality and accuracy resulted because executing requests correctly was viewed as more important for the good bureaucrat than the evil one. Both kinds of bureaucrat were expected to be more likely to grant correctly made requests than ones made incorrectly, but this effect was larger for the good bureaucrat, $M_{\text{difference}} = 0.94$, $SE = 0.1$, $p < .001$, $CI_{95\%}[0.69, 1.19]$, than the evil one, $M_{\text{difference}} = 0.52$, $SE = 0.1$, $p < .001$, $CI_{95\%}[0.25, 0.78]$. Also, the good bureaucrat was viewed as more likely than the evil one to grant requests regardless of whether they were made correctly or not. However, the effect was bigger when requests were made correctly, $M_{\text{difference}} = 0.80$, $SE = 0.1$, $p < .001$, $CI_{95\%}[0.54, 1.05]$, than incorrectly, $M_{\text{difference}} = 0.37$, $SE = 0.13$, $p = .022$, $CI_{95\%}[0.03, 0.71]$.

Finally, also consistent with Experiments 1 and 2, the interaction between understanding and accuracy resulted because although requests were viewed as more likely to be granted if recited correctly than incorrectly at both levels of understanding, the advantage for correct reciting was greater when the protagonist understood, $M_{\text{difference}} = 1.01$, $SE = 0.09$, $p < .001$, $CI_{95\%}[0.77, 1.24]$, than when the protagonist did not, $M_{\text{difference}} = 0.45$, $SE = 0.07$, $p < .001$, $CI_{95\%}[0.27, 0.64]$.

In sum, participants judgments again depended on whether requests were directed to good or evil agents. In line with the first experiment, the good bureaucrat was viewed as overall more likely than the evil one to grant requests. In line with both previous experiments, the good bureaucrat privileged requests when protagonists knew what they were requesting, whereas the evil one did not care about this. However, in contrast with the earlier experiments, the good bureaucrat cared more than the evil one about whether requests were recited correctly.

These findings suggest both similarities and differences between people's expectations about whether human and supernatural agents will grant requests. However, this conclusion is only tentative because we did not directly compare judgments about both kinds of agents. Moreover, the findings could reflect differences between vignettes used across experiments, rather than differences in expectations about supernatural beings and humans. To address this issue, the next experiment compared judgments about these two kinds of beings. Specifically, we asked participants to judge whether good and evil agents would grant various requests, while manipulating whether these agents were angels and demons (as in the first experiment) or human bureaucrats.

Also, instead of separately manipulating knowledge and accuracy as in the previous experiments, this next experiment compared three kinds of requests. In one kind of request, the

protagonist understood what they were requesting and the request was error-free. The other two requests each differed on just one of these features—the protagonist was either confused about the nature of the request *or* the request was filled with errors. This approach was necessary to avoid attempting a 4-way factorial design, as would be required if we persisted with the approach in the previous experiments, while adding the additional manipulation of whether the request-granting agent is supernatural or human.

Experiment 4

Methods

This experiment was pre-registered at <https://aspredicted.org/dq924.pdf>.

Participants. We tested 870 participants ($M_{\text{age}} = 37$ years, range = 18-80, 425 male, 440 female, and 5 indicating “other or prefer not to answer”). A further 30 participants were excluded for failing at least one of three attention check questions. Participants were recruited online using CloudResearch using the “block low quality participants” option. They were required to be U.S. residents and to have a HIT approval rate greater than 95%.

Materials & Procedure. Participants read a scenario about an agent who decides whether to approve requests in a 2X2X3 design (see Figure 4 for a sample script). We manipulated two factors-between subjects: whether the agent was good or evil, and whether the agent was a supernatural being (angel, demon) or a human bureaucrat. For all agents, instructions above the scenario read, “Regardless of your beliefs, please respond to the following story as if it is real.”

We manipulated the other factor within-subjects: Participants rated the likelihood of the agent granting three kinds of requests from a human protagonist: 1) *ideal*, the protagonist understands what she is requesting and her request is free of errors and is well written; 2)

confused, the protagonist does not understand what she is requesting, though her request is error-free and well written; 3) *errors*, the protagonist understands what she is requesting but her request is filled with errors and is poorly written.

Figure 4

Sample Script (left) and Results (right) from Experiment 4

Uriah is an angel/a demon. He resides in heaven/hell and people often write requests to him. Uriah has to decide whether to grant each request.

Uriah is exceptionally good and moral/evil and immoral. This greatly affects his decisions about whether to grant or deny requests.

A woman named Joan just wrote a request to Uriah. How likely is Uriah to grant the request if...

It's obvious that Joan understands what she is requesting. Her letter is well written and free of errors.

It's obvious that Joan is confused about what she is requesting. Her letter is well written and free of errors.

It's obvious that Joan understands what she is requesting. Her letter is poorly written and full of errors.

Note. The script shows text from the supernatural conditions; colored text was manipulated between-subjects across good and bad conditions. In conditions where the agent granting requests was a human, the top paragraph began “Uriah is a bureaucrat. He works in a municipal office...”. Error bars show 95% confidence intervals.

As in the previous experiments, each item appeared on a separate row in a matrix grid.

However, unlike the previous experiments, the grid did not include two attention check items.

Instead, the screen after the main test questions asked three multiple choice questions as attention checks. Two of these questions offered four options and one offered three, so the odds of passing all three by chance was only 1/48. We used similar attention checks in the next experiment; for questions wordings, see the scripts at <https://osf.io/njr28/>.

Results and Discussion

Figure 4 shows participants' mean ratings. We analyzed data using a model with the between-subject predictors *agent-type* (whether the agent granting the request was supernatural or human), *morality* (whether this agent was evil or good) and the within-subjects predictor *request-type* (whether the request was ideal, confused, or had errors). The three main effects were significant. Requests were overall more likely to be granted when directed to supernatural agents than human bureaucrats, Wald $\chi^2 = 61.11$, $p < .001$, and when directed to someone good rather than evil, Wald $\chi^2 = 222.45$, $p < .001$; judgments about whether requests would be granted also varied across the three request types, Wald $\chi^2 = 526.87$, $p < .001$. The three 2-way interactions were also significant: morality*request-type, Wald $\chi^2 = 234.98$, $p < .001$; morality*agent-type, Wald $\chi^2 = 43.00$, $p < .001$; granter*request-type, Wald $\chi^2 = 22.39$, $p < .001$. The 3-way interaction was non-significant, Wald $\chi^2 = 2.73$, $p = .256$.

The interaction between morality and request-type resulted because participants thought good agents would grant erroneous requests over confused requests, whereas evil agents would grant confused requests over erroneous ones. Specifically, participants thought good and evil agents would each be significantly more likely to grant ideal requests (i.e., protagonist understands and makes no errors) than confused requests (good, $M_{\text{difference}} = 1.85$, $SE = 0.06$, $p < .001$, $CI_{95\%}[1.66, 2.03]$; evil, $M_{\text{difference}} = 0.41$, $SE = 0.08$, $p < .001$, $CI_{95\%}[0.17, 0.65]$) or requests with errors (good, $M_{\text{difference}} = 0.96$, $SE = 0.06$, $p < .001$, $CI_{95\%}[0.79, 1.12]$; evil, $M_{\text{difference}} = 0.68$, $SE = 0.07$, $p < .001$, $CI_{95\%}[0.47, 0.90]$). But as the mean differences show, for good agents, the effect for confused requests is greater, whereas for evil agents it is greater for requests with errors. The cause of the interaction is especially clear when comparing the two non-ideal requests. Participants thought good agents would be more likely to grant requests with errors

than confused requests, $M_{\text{difference}} = 0.89$, $SE = 0.06$, $p < .001$, $CI_{95\%}[0.70, 1.08]$. However, participants thought evil agents would show the reverse pattern and would be more likely to grant requests with errors than confused requests, $M_{\text{difference}} = -0.27$, $SE = 0.06$, $p < .001$, $CI_{95\%}[-0.45, -0.09]$.

Examining this interaction further, participants also thought good agents would be more likely than evil agents to grant requests whenever the protagonist understood what she was requesting (ideal, $M_{\text{difference}} = 1.32$, $SE = 0.07$, $p < .001$, $CI_{95\%}[1.12, 1.53]$; errors, $M_{\text{difference}} = 1.05$, $SE = 0.08$, $p < .001$, $CI_{95\%}[0.83, 1.28]$). But expectations about good and evil agents did not significantly differ when protagonists were confused about what they were requesting, $M_{\text{difference}} = -0.11$, $SE = 0.08$, $p > .999$, $CI_{95\%}[-0.33, 0.12]$.

The other 2-way interactions are less central for our main questions. Nonetheless, the interaction between morality and agent-type resulted because participants thought requests were more likely to be granted by demons than evil bureaucrats, $M_{\text{difference}} = 0.73$, $SE = 0.79$, $p < .001$, $CI_{95\%}[0.52, 0.94]$, but they did not have corresponding expectations for good agents (i.e., ratings for angels and good bureaucrats did not significantly differ), $M_{\text{difference}} = .06$, $SE = .06$, $p > .999$, $CI_{95\%}[-0.10, 0.23]$. Finally, the interaction between agent- and request-type resulted because participants thought supernatural agents would be more likely than humans to grant both kinds of non-ideal requests (confused, $M_{\text{difference}} = 0.46$, $SE = 0.08$, $p < .001$, $CI_{95\%}[0.23, 0.68]$; errors, $M_{\text{difference}} = 0.58$, $SE = .08$, $p < .001$, $CI_{95\%}[0.36, 0.81]$), whereas their ratings between these two kinds of agents did not significantly differ for ideal requests, $M_{\text{difference}} = 0.15$, $SE = 0.07$, $p = .484$, $CI_{95\%}[-.056, 0.36]$.

In sum, participants thought good agents would be more sensitive to whether the protagonist was confused than to whether her request had errors. In contrast, participants thought

evil agents would be more sensitive to errors in the request than to whether the protagonist was confused. Although we found some differences between judgments about supernatural and human agents, these differences did not impact the interaction between morality and request-type (i.e., there was no 3-way interaction). In contrast with the previous experiments, participants in this experiment indicated that evil agents would show *some* sensitivity to whether the protagonist understood the request. One reason for this difference may be that this experiment used a very large sample, increasing the possibility of detecting small effects. This difference could also reflect changes in the scenarios across stories (we revisit this idea in the General Discussion).

The results so far show that people expect evil agents to be relatively insensitive to intentions when granting requests. But why do people expect this? One possibility is that people expect evil agents to be inattentive. For example, if the evil agent is careless when reading Joan's letter, he might not even notice her confusion. A second possibility is that people think evil agents are incapable of understanding others' intentions; people might conceive of evil agents as lacking empathy and the ability to take on others' perspectives. The final possibility we consider is that evil agents are viewed as indifferent to others' intentions. For example, the demon might fully realize that Joan is confused about what she is requesting but might grant the request anyway because he gives her intentions little weight.

To test between these accounts, we had participants read an updated version of the scenarios where Joan writes to an angel or demon. In this version, Joan was confused about what she was requesting, and the supernatural agent grants her request. Participants then rated explanations for why he granted her request. This design allows us to examine which explanations distinguish participants' conceptions of evil agents from their conceptions of good

ones. As pre-registered (see below), we thought the key factor would be the expectation that evil agents are more likely than good ones to be indifferent to others' intentions.

Experiment 5

Methods

This experiment was pre-registered at <https://aspredicted.org/yn3dd.pdf>.

Participants. We tested 439 participants ($M_{\text{age}} = 40$ years, range = 18-92, 236 male, 201 female, and 2 indicating “other or prefer not to answer”). A further 14 participants were excluded for failing at least one of three attention check questions. Participants were recruited online using CloudResearch using the criteria from the previous experiment but barring participants who had completed it.²

Materials & Procedure. Participants read a scenario about an angel or demon (manipulated between-subjects) who decides to approve a request from a woman who is confused about what she is requesting (see Figure 5 for the script); text above the scenario again asked participants to respond as though the scenario was real. Each participant then rated four explanations for why the agent granted this request. Participants rated how likely each explanation was using a 5-point Likert scale ranging from *very unlikely* (1) to *very likely* (5). As in the previous experiments, the items appeared together in a matrix grid.

The first three explanations were that the agent was *inattentive* and did notice the woman was confused, *incapable* of understanding that a person could be confused about their request, and *indifferent* to the woman being confused about her request. The presentation order of these explanations was randomized across participants. The final explanation (i.e., always shown last) was that the agent had *other reasons* for granting the request. We mainly included this

² We did not take similar precautions in the earlier experiments, and so one limitation of the present work is that some participants could have contributed data to more than one experiment.

explanation so that participants could give one item a high rating, even if the rest seemed implausible. This was particularly a concern for good agents, who might reasonably be viewed as attentive, capable, and concerned about what Joan meant.

Figure 5

Script (left) and Results (right) from Experiment 5

Uriah is an angel/a demon. He resides in heaven/hell and people often write requests to him. Uriah has to decide whether to grant each request.

A woman named Joan wrote a request to Uriah. It was obvious that Joan was confused about what she requested. After reading her letter, Uriah decided to grant the request.

Below are four explanations for why Uriah granted Joan's request even though she was confused. Please rate how likely each is.

Uriah didn't read the letter carefully, so he didn't notice that Joan was confused.

Uriah isn't capable of understanding when people are confused about what they are requesting.

Uriah knew that Joan was confused, but he didn't care about what she really intended.

Uriah had other reasons to grant Joan's request.

Note. Colored text was manipulated between-subjects. Error bars show 95% confidence intervals.

Results and Discussion

Figure 5 shows participants' mean ratings. We analyzed data using a model with the between-subject predictor *moral character* (good, evil) and the within-subjects predictor *explanation* (inattentive, incapable, indifferent, other reasons). This interaction revealed a main effect of morality, Wald $\chi^2 = 53.38$, $p < .001$, a main effect of explanation, Wald $\chi^2 = 1059.36$, $p < .001$, and a significant interaction between these factors, Wald $\chi^2 = 184.12$, $p < .001$.

The interaction resulted because participants more strongly endorsed the indifferent explanation when the agent was a demon than an angel, $M_{\text{difference}} = 1.67$, $SE(0.11)$, $p < .001$, $CI_{95\%}[1.34, 2.01]$, whereas moral character did not significantly affect ratings for other explanations: inattentive, $M_{\text{difference}} = -0.05$, $SE(0.10)$, $p > .999$, $CI_{95\%}[-0.38, 0.27]$; incapable, $M_{\text{difference}} = 0.04$, $SE(0.12)$, $p > .999$, $CI_{95\%} [-0.34, 0.41]$; other reasons, $M_{\text{difference}} = 0.07$, $SE(0.1)$, $p > .999$, $CI_{95\%} [-0.24, 0.37]$.

Also, participants endorsed indifference more strongly than inattentiveness or incapability regardless of the agent's moral character. However, the effects were greater for evil agents (inattentive, $M_{\text{difference}} = 2.55$, $SE(0.10)$, $p < .001$, $CI_{95\%}[2.24, 2.86]$; incapable, $M_{\text{difference}} = 2.29$, $SE(0.11)$, $p < .001$, $CI_{95\%}[1.94, 2.63]$) than good ones (inattentive, $M_{\text{difference}} = 0.82$, $SE(0.11)$, $p < .001$, $CI_{95\%}[0.49, 1.15]$; incapable, $M_{\text{difference}} = 0.65$, $SE(0.11)$, $p < .001$, $CI_{95\%}[0.31, 0.98]$). Finally, when the agent was an angel, participants thought having other reasons for granting the request was a better explanation than indifference, $M_{\text{difference}} = 1.35$, $SE(0.12)$, $p < .001$, $CI_{95\%} [0.98, 1.73]$. But when the agent was a demon, ratings of these explanations did not significantly differ, $M_{\text{difference}} = 0.25$, $SE(0.08)$, $p = 0.073$, $CI_{95\%}[-0.01, 0.52]$.

In sum, indifference was more strongly endorsed when the agent was evil rather than good. Participants did not think that inattentiveness or incapability was a good explanation, regardless of the agent's morality. We also found that participants often thought both kinds of agents might have other reasons for granting confused requests.

General Discussion

Our findings suggest that people expect morally good and bad agents to differ in their sensitivity to the intentions behind requests. Specifically, participants thought that morally good

agents would be far more likely than evil agents to attend to the intentions behind a request when deciding whether to grant it. This pattern was consistent regardless of whether participants considered requests to supernatural agents or humans, and was also robust regardless of whether requests were made by reading a passage aloud, filling out paperwork, or writing a letter. These findings suggest that people have distinct expectations of how moral character affects decision-making, and also suggest that people's beliefs about how good and evil supernatural agents will respond to requests are grounded on their beliefs about how humans do so.

We also found that moral character affected participants' expectations about agents' sensitivity to whether requests were executed correctly or with errors. One consistent pattern here was that participants expected evil agents to be more influenced by the requesters' accuracy than intentions, while expecting the opposite from good agents. But besides this, the findings varied somewhat across the experiments. In some, participants expected evil agents to be more sensitive than good agents to accurate execution, whereas other experiments did not find this. Although further work is needed, we suspect these differences could pertain to background beliefs about the various scenarios we used. For example, familiarity with religious and fictional depictions of demons might lead people to expect them to be more likely to respond to recited incantations than to personally written letters. So, the method used to convey requests might have constrained the extent to which ratings could depend on the protagonist's accuracy.

We also examined beliefs about *why* evil agents would grant confused requests, which do not reflect requesters' actual intentions. In advance, it seemed these beliefs could pertain to agents' awareness of others' intentions. For example, whereas people might expect good agents to notice the intentions behind requests, they might expect evil agents to carelessly overlook these intentions. Or people might even believe that evil agents lack the ability to understand that

some requests are not backed by intentions—as might be expected if people view evil agents as oblivious to others’ mental states. However, participants did not strongly endorse these awareness-based explanations. Instead, they believed evil agents would recognize when requests are not backed by intentions but thought evil agents would not care. So people viewed indifference to intentions, rather than ignorance of them, as the better explanation for why evil agents will grant confused requests.

Our findings build on previous work showing that people view traits like trustworthiness, fairness, and loyalty to be informative about moral character, and that these traits strongly influence overall impressions of others (e.g., Goodwin et al., 2014). Specifically, our findings suggest that people have distinct models of how agents with good and bad moral character make decisions, and they also suggest these models differ in the weight given to information about others’ intentions. This said, much remains to be discovered. One possibility is that good agents are thought to be particularly sensitive to intentions, but another is that evil agents are thought to be particularly likely to disregard them. Or perhaps people hold both of these views. Testing between these possibilities might require asking participants about morally neutral agents (e.g., Jinn are often viewed this way, but might be viewed as not having the volition to decline to perform requests). For example, if participants thought such agents would be similar to good ones in considering the intentions behind requests, this might suggest our findings were mostly driven by beliefs about evil moral character.

There is also much more to learn about other ways models of good and evil agents differ. Our findings suggest people do not hold straightforward expectations of good agents acting to help others, and evil agents acting to inflict harm and suffering (Baumeister, 1999; Webster & Saucier, 2013). For example, in the first experiment, participants did not uniformly predict that

angels would grant the positive request to protect the protagonist's love. Instead, these predictions varied based on whether the protagonist understood the request and uttered it correctly. In advance, it might also have seemed that people should believe that good agents are overall more likely to grant requests than evil agents, or that evil agents would maliciously deny all requests. But again, our findings suggest otherwise. For example, participants were more likely to predict that requests would be granted by the devil than by god. Much of this variation in participants' judgments may have stemmed from differences in participants' background assumptions about the various scenarios. For example, participants might have assumed the devil would benefit by granting requests, but they might not have assumed this about evil bureaucrats.

Our findings may also relate to work demonstrating other ways that manipulations of good versus bad affect people's judgments (e.g., Knobe, 2010; Phillips et al., 2014; Sussman & Oppenheimer, 2020; Newman et al., 2014). Our findings may be especially reminiscent of work showing that people often view morally bad outcomes as more intentional than morally good ones (e.g., Knobe, 2003a, 2003b). For example, if the owner of factory initiates a new project to increase profits while foreseeing it will harm the environment, people view the harm as intentional. But if the owner foresees that the project will benefit the environment, people conclude the resulting benefits are unintentional. One potential overlap between such findings and the present work is that both fit with the idea that immoral agents will often be indifferent to considerations outside their main aims and focus—harming the environment in the story about the factory owner, and ignoring the requester's intentions in our studies (see Uhlmann et al., 2015 for related discussion).

This picture of evil agents as indifferent to considerations outside their main focus might help explain why participants expected them to prioritize how requests were executed over the

intentions behind them. When an agent decides whether to grant a request, errors in its execution might be viewed as falling within the agent's focus, whereas intentions behind the request might be viewed as an outside consideration. This idea could be further investigated by examining how other factors affect predictions about whether good and evil agents will grant requests. For example, future work could specify whether fulfilling requests will harm or benefit the individual making the request (or other people too), whether this person is themselves good or evil, and whether fulfilling the request will benefit the agent who grants it. We might find, for example, that evil agents will be expected to fulfill requests that benefit others, so long as this opportunistically serves the agents' own aims. In contrast, we might not expect good agents to only grant requests that have good outcomes, or that come from morally good individuals. Exploring these various possibilities could help elucidate what participants had in mind when they strongly endorsed "other reasons" as explaining why both good and evil agents would grant confused requests.

Finally, although our findings suggest people believe that evil agents choose to ignore others' intentions when granting requests, we do not yet know how far this belief generalizes. At one extreme, people might only expect evil agents to show indifference to intentions when deciding whether to grant requests. However, if people view evil agents as generally indifferent to considerations outside their main aims and focus, evil agents might be expected to show this indifference in many other contexts. For example, whereas individuals are usually viewed as more blameworthy if they cause harm intentionally rather than unintentionally, people might expect evil agents to be indifferent to intentions when blaming others. This view of evil agents also predicts that there will be circumstances where they will give ample attention to other's intentions. Cruelty and malice often hinge on humiliating others and inducing shame and other

negative emotions (Bloom, 2017). Hence, when an evil agent is set on harming someone, they might be expected to give extra attention to that person's intentions—the better to frustrate them.

Conclusions

When judging whether both supernatural agents and humans would grant requests, participants expected greater sensitivity to the intentions behind requests from good agents than evil ones. When endorsing explanations for why an agent would grant a confused request, participants more strongly endorsed indifference when the agent was evil. Together, these findings suggest people have distinct models of how agents with good and bad moral character weigh others' intentions when making decisions, and provide some support for the broader claim that evil agents are viewed as indifferent to considerations that do not directly pertain to their own aims.

Open Practices. The full data and all scenario scripts are available at <https://osf.io/njr28/>. Pre-registrations are available for Experiment 2, <https://aspredicted.org/c9xb4.pdf>; Experiment 3, <https://aspredicted.org/a8u3y.pdf>; Experiment 4, <https://aspredicted.org/dq924.pdf>; and Experiment 5, <https://aspredicted.org/yn3dd.pdf>.

References

- Barrett, J. (2002). Smart gods, dumb gods, and the role of social cognition in structuring ritual intuitions. *Journal of Cognition and Culture*, 2(3), 183–193.
<https://doi.org/10.1163/15685370260225080>
- Barrett, J. L., & Keil, F. C. (1996). Conceptualizing a nonnatural entity: anthropomorphism in God concepts. *Cognitive Psychology*, 31(3), 219–247.
<https://doi.org/10.1006/cogp.1996.0017>
- Baumeister, R. F. (1996). *Evil: Inside human cruelty and violence*. W H Freeman/Times Books/Henry Holt & Co.
- Benkovitz, M. J. (1950). The good serpent. *Folklore*, 61(3), 146-151.
<https://doi.org/10.1080/0015587X.1950.9718001>
- Bloom, P. (2017). The root of all cruelty. *The New Yorker*, 74-77. Retrieved from
<https://www.newyorker.com/magazine/2017/11/27/the-root-of-all-cruelty>
- Epley, N., Morewedge, C. K., & Keysar, B. (2004). Perspective taking in children and adults: Equivalent egocentrism but differential correction. *Journal of Experimental Social Psychology*, 40(6), 760-768. <https://doi.org/10.1016/j.jesp.2004.02.002>
- Garrard, E. (2002). Evil as an explanatory concept. *The Monist*, 85(2), 320-336.
<https://doi.org/10.5840/monist200285219>
- Goodwin, G. P. (2015). Moral character in person perception. *Current Directions in Psychological Science*, 24(1), 38-44. <https://doi.org/10.1177/0963721414550709>
- Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, 106(1), 148–168. <https://doi.org/10.1037/a0034726>

- Gromet, D. M., Goodwin, G. P., & Goodman, R. A. (2016). Pleasure from another's pain: The influence of a target's hedonic states on attributions of immorality and evil. *Personality and Social Psychology Bulletin*, 42(8), 1077-1091.
<https://doi.org/10.1177%2F0146167216651408>
- Heiphetz, L., Lane, J. D., Waytz, A., & Young, L. L. (2016). How children and adults represent God's mind. *Cognitive Science*, 40(1), 121-144. <https://doi.org/10.1111/cogs.12232>
- Heiphetz, L., Lane, J. D., Waytz, A., & Young, L. L. (2018). My mind, your mind, and God's mind: How children and adults conceive of different agents' moral beliefs. *British Journal of Developmental Psychology*, 36(3), 467-481.
<https://doi.org/10.1111/bjdp.12231>
- Helzer, E. G., Critcher, C. R. (2018). What do we evaluate when we evaluate moral character? In K. Gray & J. Graham (Eds.), *Atlas of moral psychology* (pp. 99–107). Guilford Press.
- Hepach, R., Benziad, L., & Tomasello, M. (2020). Chimpanzees help others with what they want; children help them with what they need. *Developmental Science*, 23(3), e12922.
<https://doi.org/10.1111/desc.12922>
- Knobe, J. (2003a). Intentional action and side effects in ordinary language. *Analysis*, 63(3), 190-194.
- Knobe, J. (2003b). Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical psychology*, 16(2), 309-324.
<https://doi.org/10.1080/09515080307771>
- Knobe, J. (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences*, 33, 315-329. <https://doi.org/10.1017/S0140525X10000907>

- Landy, J. F., Piazza, J., & Goodwin, G. P. (2016). When it's bad to be friendly and smart: The desirability of sociability and competence depends on morality. *Personality and social psychology bulletin*, 42(9), 1272-1290.
- Lane, J. D., & Dolins, F. (2016). Socio-cultural differences in judgments about the power of thought. In A. Village & R. W. Hood (Eds.), *Research in the social scientific study of religion*, Volume 27 (pp. 174-191). Brill. https://doi.org/10.1163/9789004322035_012
- Lane, J. D., Evans, E. M., Brink, K. A., & Wellman, H. M. (2016). Developing concepts of ordinary and extraordinary communication. *Developmental Psychology*, 52(1), 19-30. <https://doi.org/10.1037/dev0000061>
- Martin, A., & Olson, K. R. (2013). When kids know better: Paternalistic helping in 3-year-old children. *Developmental Psychology*, 49(11), 2071-2081. <https://doi.org/10.1037/a0031715>
- Martyrdrevel27 (2013, May 24). *I think I might have accidentally sold my soul, advice?* [Online forum post]. EscapistMagazine.com. <https://forums.escapistmagazine.com/threads/i-think-i-might-have-accidentally-sold-my-soul-advice.118460/>
- Nadler, J. (2012). Blaming as a social process: The influence of character and moral emotion on blame. *Law and contemporary problems*, 75(2), 1-31.
- Newman, G. E., Bloom, P., & Knobe, J. (2014). Value judgments and the true self. *Personality and Social Psychology Bulletin*, 40(2), 203-216. <https://doi.org/10.1177/0146167213508791>
- Pettit, P. (2018). Three mistakes about doing good (and bad). *Journal of Applied Philosophy*, 35(1), 1-25. <https://doi.org/10.1111/japp.12294>

- Phillips, J., Nyholm, S., & Liao, S. Y. (2014). The good in happiness. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy: Volume 1* (pp. 253-293). Oxford University Press.
- Siegel, J. Z., Crockett, M. J., & Dolan, R. J. (2017). Inferences about moral character moderate the impact of consequences on blame and praise. *Cognition*, 167, 201–211.
<https://doi.org/10.1016/j.cognition.2017.05.004>
- Shaman, N. J., Saide, A. R., Lesage, K. A., & Richert, R. A. (2016). Who cares if I stand on my head when I pray? Ritual inflexibility and mental-state understanding in preschoolers. In A. Village & R. W. Hood (Eds.), *Research in the social scientific study of religion, Volume 27* (pp. 122-139). Brill.
https://doi.org/10.1163/9789004322035_009
- Shtulman, A., & Lindeman, M. (2016). Attributes of God: Conceptual foundations of a foundational belief. *Cognitive Science*, 40(3), 635-670.
<https://doi.org/10.1111/cogs.12253>
- Shtulman, A., & Rattner, M. (2018). Theories of God: Explanatory coherence in religious cognition. *PloS one*, 13(12), e0209758. <https://doi.org/10.1371/journal.pone.0209758>
- Sussman, A. B., & Oppenheimer, D. M. (2020). The effect of effects on effectiveness: A boon-bane asymmetry. *Cognition*, 199, 104240.
<https://doi.org/10.1016/j.cognition.2020.104240>
- Takács, T. (Director) (1987). The gate [Film]. New Century Entertainment; Alliance Entertainment.
- u/12212888 (2020, January 9). *Im worried i might have accidentally sold my soul?* [Online forum post]. Reddit.

https://www.reddit.com/r/Christianity/comments/emlqgl/im_worried_i_might_have_accidentally_sold_my_soul/

Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A person-centered approach to moral judgment. *Perspectives on Psychological Science*, 10(1), 72-81.

Webster, R. J., & Saucier, D. A. (2013). Angels and demons are among us: Assessing individual differences in belief in pure evil and belief in pure good. *Personality and Social Psychology Bulletin*, 39(11), 1455-1470. <https://doi.org/10.1177%2F0146167213496282>

Young, L., & Tsoi, L. (2013). When mental states matter, when they don't, and what that means for morality. *Social and Personality Psychology Compass*, 7(8), 585-604.
<https://doi.org/10.1111/spc3.12044>