

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/313077052>

# The uncensored truth about morality

Article · December 2015

CITATIONS  
10

READS  
218

3 authors:



**Chelsea Schein**

University of North Carolina at Chapel Hill

17 PUBLICATIONS 419 CITATIONS

SEE PROFILE



**Amelia Goranson**

University of North Carolina at Chapel Hill

2 PUBLICATIONS 29 CITATIONS

SEE PROFILE



**Kurt Gray**

University of North Carolina at Chapel Hill

100 PUBLICATIONS 3,408 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Technology and Sociability [View project](#)



Uncanny Valley [View project](#)

Running Head: Morality

Schein, C., Goranson, A., & Gray, K. (2015). The uncensored truth about morality. *The Psychologist*, 28(12), 982-985.

The uncensored truth about morality

Chelsea Schein, Amelia Goranson & Kurt Gray

University of North Carolina, Chapel Hill

Word Count: 3,255

Corresponding author:

Kurt Gray  
Department of Psychology  
University of North Carolina, Chapel Hill  
Chapel Hill, NC 27599

Email: [kurtgray@unc.edu](mailto:kurtgray@unc.edu)

## Abstract

Morality is about many things, but is ultimately about harm. All people appear to understand morality through the same harm-based template of two perceived minds—a thinking doer who causes suffering to a vulnerable feeler. This “dyadic moral template” aligns with most obvious cases of evil (e.g., murder, child abuse), and also explains why moral disagreement centers on issues in which harm is ambiguous (e.g., prostitution, pornography). We provide an introduction to this new theory, framing it in the context of recent debates about pornography and censorship.

The uncensored truth about morality: Ubiquitous harm and our dyadic moral minds

In December 2014, the British Board of Film Censors (BBFC) banned pornographic materials depicting acts such as consensual spanking, facesitting, urolagnia, and female ejaculation. According to a spokesman for the organization, the “BBFC classification regime is a tried and tested system of what content is regarded as harmful for minors” and the bill was aimed at “safeguarding children” (Quinn, 2014). However, critics of the bill quickly pointed out that this list is not only “arbitrary,” but includes only harmless acts between consenting adults. The BBFC and its critics seem to see the world in two very different ways—one in which pornography is *both* immoral *and* harmful to children, and one in which pornography is *neither* immoral *nor* harmful. This alignment of perceived immorality and perceived harm is not coincidental, but instead reflects the central role of harm in moral cognition.

The power of harm within morality is intuitive, as many of the gravest moral violations are those that directly cause harm; murder, assault, theft, rape, and abuse are acts that are both harmful and universally viewed as immoral. Harm is also central to rhetoric about immorality, as the issue of censorship clearly illustrates. When people want to argue that an act is wrong, they reference its harmfulness, often to children. However, some recent theories of morality (e.g., moral foundations theory; Haidt, 2012) have suggested that this harm-based rhetoric is just that—mere rhetoric. These theories claim that questions of harm are relevant to only a fraction of moral judgments, and that arguments about harm reflect only convenient post-hoc justifications (Haidt, 2012).

However, the new theory of *dyadic morality* pulls harm from the shadows back into the spotlight (Gray, Waytz, & Young, 2012). Dyadic morality suggests that harm is not only the most important factor in moral judgment, but the very core of a universal moral template—a template which both detects harm in diverse moral transgressions (Schein & Gray, in press) and

fills it in when apparently absent (Gray, Schein, & Ward, 2014). In this article, we first describe dyadic morality before addressing three apparent challenges to this theory: the intuitive nature of morality (intuitionism), the moralization of “harmless” transgressions in other cultures (pluralism), and cognitive differences across moral content (modularity).

### *Categorization and the Moral Dyad*

In essence, moral judgment is simply about categorization. The question “is X immoral?” can be reworded as “does X belong to the category of immorality?” As decades of research in cognitive psychology suggest, stimuli are categorized by automatically comparing them to a category’s prototype (Murphy, 2004). This prototype (or “cognitive template”) emerges from the most common, salient, and important features of the category. For example, the template of “bird” is something small, winged, seed eating, and capable of flight. The better a stimulus matches the template, the more robustly it is categorized as part of that group, explaining why sparrows are judged to be birds faster than ostriches.

What is the template of immorality? It is based upon acts that are the most universally condemned, the most evolutionary important, and the most emotionally evocative. Harm fulfills all these criteria. Harmful acts such as murder and assault are condemned across cultures, strongly impair genetic survival, and consistently evoke powerful negative emotion (Baron-Cohen, 2011). Of course, many harmful acts (e.g., car accidents) are not immoral, so our moral template must have some additional features.

Research finds that our moral template involves two interacting minds—an intentional agent (i.e., perpetrator) harming a suffering patient (i.e., victim; Gray et al., 2012). In other words, the essence of immorality isn’t merely “harm” but “harm caused by an agent.” This agent can be another person, a corporation, or a government, but it must be an entity perceived to be

mentally capable of intention and action. Conversely, the patient who receives the harm must be mentally capable of feeling pain and suffering, such as children, or puppies or the elderly.

The combination of “intentional agent and suffering patient”—or even more simply “thinking doer and vulnerable feeler”—gives us a simple formula for understanding the moral world. People are most morally incensed when powerful thinking doers harm powerless vulnerable feelers (e.g., a CEO kicks a baby). This is exactly the combination of minds seen by the BBFC censors, who believed that profit-driven movie producers were harming children. Conversely, people are *least* morally incensed when powerless vulnerable feelers harm powerful thinking doers (e.g., a baby bites a CEO).

With this formula of “thinking doer and vulnerable feeler,” we can predict people’s moral outrage to any infraction with two questions: How much does the patient/victim suffer? How much does the agent/perpetrator intend the harm? We can also predict a situation’s *potential* for moral outrage by considering the patient’s potential for suffering and the agent’s potential for thought and action.

Because this moral template involves two interacting minds, it is called the *dyadic* template, from the Greek work *dyo*, meaning two. Unlike the template for birds (or dogs or furniture), the moral template is a matter of *perception*. Whether a bird can fly is a matter of fact, but less certain is whether a perpetrator is capable of intentional thought, or a victim is capable of suffering. The inherent ambiguity of other minds means that good people can nevertheless have moral disagreement, because they see different mental capacities in potential perpetrators and victims. Debates about whether children who kill should go to adult prison hinge on whether they are fully capable thought and action—an issue of perception as much as fact. Likewise, debates about abortion hinge on whether fetuses are capable of feeling pain—

which is also largely a matter of perception. Where people perceive minds, they perceive the potential for evil.

### *Moral Universals*

The dyadic template not only allows for moral differences between people, but also suggests two powerful moral universals. The first is *what is harmful seems wrong* (Schein & Gray, in press). As we mentioned above, stimuli (e.g., sparrows) are judged as members of a category (e.g., birds) based on their similarity to a template. As the moral template is rooted in perceptions of harm, the more apparently harmful an act, the more it is judged as immoral. This explains why first-degree murder (obviously intentional and harmful) is universally condemned, whereas pornography (ambiguous intention and harm) is less universal condemned—and hence a matter of debate. More technically, we can say that harm is central to moral judgment; it is the key criterion for determining whether or not something is immoral.

Our lab tested the centrality of harm in a number of studies (Schein & Gray, in press). One asked participants to volunteer the first immoral act that came to mind. If harm is central to moral judgment, people should volunteer something obviously harmful, which is exactly what they did. Over 90% of the acts recalled were dyadic in nature, such as murder, abuse, theft and adultery. Another test asked people to rate the immorality of acts that violated different kinds of norms, for example those that were harmful, unfair, disloyal, disrespectful, or gross. Consistent with a dyadic template, harmful acts were seen as the most immoral. The importance of harm was also observed in automatic judgments, as reaction times to categorize an action as “immoral” almost perfectly predicted reactions times to categorize an action as “harmful.” Whether an act is immoral seems best predicted by its perceived harmfulness.

The second moral universal predicted by dyadic morality is *what is wrong seems harmful* (Gray et al., 2014). Imagine you knew nothing about an action beyond that people believed it to be truly evil. You would assume that this action was harmful—or at least that these people *perceived* it to be harmful. This perceived harm is an inevitable consequence of a dyadic moral template, which not only determines how stimuli are categorized, but also shapes how stimuli are perceived *after* categorization. This is because they are not static psychological structures, but lenses (or schemas) through which we view the world. Like any lens, they powerfully shape our perceptions, often without our knowledge.

Consider again the template of birds. Because the concept of flying is central to birds, you will automatically assume that something labeled “bird” can fly. Because the concept of “harm” is central to immorality, you will automatically assume that something labeled “immoral” is harmful. Of course, with birds, you can objectively learn that some birds cannot fly. But because harm is subjective, people can seldom “objectively” learn that something they see as evil is also harmless. Returning to the example of censorship, no amount of cajoling about the harmlessness of consensual spanking or facesitting is going to change perceptions of those who already perceive it as harmful (and wrong). No one ever says “It’s harmless and wrong!”

This link from wrong to harm (and back) isn’t just rhetoric, but automatic and intuitive. In one study from our lab, people read about acts that seemed wrong despite causing no directly physical harm, such as masturbating to a picture of your dead sister (yes, we know it’s bizarre—blame our reviewers). Consistent with a dyadic template, participants nevertheless labeled these “objectively harmless” misdeeds as harmful. Were these perceptions mere effortful justification? No. In fact, people were *especially* likely to see harm when we impaired their ability to reason by forcing them to answer quickly.



Other studies reveal that these perceptions of harm even creep into other judgments. After reading about “harmless immoral acts,” people were more likely to see sad expressions in the faces of children—an experimental demonstration of why those against pornography inevitably see it as harming children. Importantly, perceptions of harm weren’t driven by general feelings of “badness,” as other judgments about these children were not more negative. Perceived harm leads to judgments of evil, and judgments of evil lead *specifically* to perceived harm.

### *Questions for Dyadic Morality*

Although dyadic morality aligns with decades of research on categorization and historical harm-centric accounts of moral judgment (Turiel, 1983), it conflicts with one popular theory of morality—moral foundations theory (MFT; Haidt, 2012). Moral foundations theory argues against the overarching role of harm, instead advocating for distinct moral modules, each corresponding to different kinds of moral content, such as “purity” or “fairness”. In these theories, harm is merely one “little switch” of the moral mind, and not a global template. Advocates of MFT have challenged dyadic morality primarily along three lines: intuitionism, moral pluralism, and content differences. We summarize and address these challenges here, using them as an opportunity to correct common misconceptions about the theory of dyadic morality.

#### *Intuitionism*

Historic accounts of moral judgment emphasized the importance of harm, and also emphasized the importance of moral reasoning. These theories—such as Kohlberg’s stage model (1969)—suggested that moral judgment relied upon careful deliberation and conscious reflection. Reacting against this “rule of reason,” more recent accounts have claimed that

morality is typically a matter of intuitive judgments and affective reactions (Haidt, 2001). When people denounce (or fail to denounce) pornography, it is rarely because they have rationally considered the issue. Instead, their moral judgments reflects emotion based intuitions. As one analogy suggests, our faculty of moral judgment is not an impartial judge weighing the evidence, but rather, it is an impassioned lawyer arguing in favor of its original position (Haidt, 2012).

Significant evidence supports the power of intuition in determining our moral judgment. Moral judgments are formed quickly, are influenced by incidental emotions, and can sometimes seem contradictory (Greene, 2013). As dyadic morality embraces the power of harm, some have assumed that it also embraces the reign of reason. Nothing could be further from the truth. The role of templates in categorization—whether moral or non-moral—progresses intuitively and automatically. When people think about birds, they need not consciously ponder about the importance of flight as they categorize sparrows and ostriches. Instead, key template features are used automatically and effortlessly in categorization.

As we reviewed above, studies reveal that the role of harm in moral judgment is both intuitive and automatic—consistent with a dyadic moral template. People rapidly see harm in moral violations, automatically use this perceived harm in forming their moral judgments, and effortlessly perceive harm in response to immorality (Gray et al., 2014; Schein & Gray, in press). Of course, people can (and often do) use considerations of harm in subsequent deliberative moral reasoning , but initial perceptions of harm are automatic.

### *Moral Pluralism*

Anthropologists have long recognized that different cultures have different morals. In the West, there are no clear prohibitions about what you can eat after a loved one's death. On the other hand, in India, Oriya Hindu Brahmans believe it to be immoral for the eldest son to eat

chicken after his father's death (Shweder, 2012). To account for these differences, theories such as MFT suggest that our basis for morality must extend beyond harm to considerations of spiritual purity (Haidt, 2012). This perspective is called moral pluralism, because it advocates for a plurality of moral concerns. On the surface, moral pluralism seems to argue against the harm-centric dyadic morality, however, a closer inspection suggest that dyadic morality is actually *more* consistent with moral pluralism than is MFT.

Although MFT embraces non-Western conceptions of morality, it fails to embrace non-Western notions of harm, rigidly defining it as only direct physical suffering. In contrast, dyadic morality suggests that harm is in the eye of the beholder, which means that even violations of spiritual purity can actually be grounded in concerns about harm. Consider the example of the Brahman son eating chicken—Hindus believe that the eldest son is responsible for processing the father's "death pollution" by eating a vegetarian diet (Shweder, 2012). By eating chicken, the son is thereby condemning the father's soul to eternal suffering. In other words, what superficially seems to just about "purity," is actually about harm. This idea is also supported by research from our lab, which finds that violations of "purity" are judged to be simply a certain kind of perceived harm (Gray & Keeney, 2015).

Our research also finds little support for another oft-discussed "cultural" difference in morality—MFT has suggested that US liberals and conservatives have different "foundations" (Haidt, 2012). While it is true that liberals and conservatives often morally disagree, they share similar harm-based moral judgments, as revealed by studies in our lab. Conservatives oppose gay marriage not because it seems "impure," but because it seems harmful. One notable anti-gay activist believed that allowing gay rights would destroy the American family, throw society into chaos and—of course—harm children. As with the Hindu Brahmans, MFT dismisses these

perceptions of harm as illegitimate, whereas dyadic morality embraces “harm pluralism”—the diversity of perceived harm. With its embrace of both moral and harm pluralism, dyadic morality better embodies the anthropological tradition that MFT seeks claim for itself.

### *Content Differences*

Dyadic morality suggests that—consistent with other forms of judgment—moral judgment is based upon comparison to an overarching template. In contrast, MFT believes that moral judgment is underlain by cognitive modules or “little switches in the brain” (Haidt, 2012, p. 123), each of which have fundamentally “distinct cognitive computations” (Young & Saxe, 2011, p. 203). This module view suggests that violations of one “content” area (e.g., harmful acts like murder) are processed differently from those of another (e.g., purity acts like facesitting). The module view also suggests that different kinds of moral content are linked to different emotions, such that harm (e.g., murder) is linked to anger, and purity (e.g., facesitting) is linked to disgust.

Despite the common acceptance of these claims in our field (Haidt, 2012), our lab wondered if all was not as it seemed. We examined the research arguing for unique cognition and unique emotions in the content area of “purity,” and found no support for either claim. First, uniqueness requires distinctness, and judgments of “harm” and “purity” are highly correlated ( $r = .87$ ; Gray & Keeney, 2015). Second, previous studies had clear confounds—unaccounted for 3rd variables—which we suspected gave only the illusion that purity was special.

Studies arguing for unique cognition all used a specific set of scenarios. These scenarios represented harm through acts such as murder and child abuse, and represented purity through acts such as masturbating with a dead chicken or getting a tail via plastic surgery. The savvy reader will recognize that these scenarios differ not only in moral “content” but also in severity

and weirdness. The harm scenarios are punishable by decades in prison, unlike the purity scenarios. The harm scenarios are also much less *weird* (i.e., atypical) than the purity scenarios. We read about murder everyday in the paper, and can easily imagine motivations for killing another person. Conversely, we seldom read about tail-plastic-surgery, and have difficulty imagining why someone would do this. Our studies reveal that these differences in severity and weirdness account for the apparent uniqueness of purity. With careful experimental controls, the specialness of purity disappears altogether (Gray & Keeney, 2015).

Careful experimental controls also cause the apparent link between purity and disgust to vanish. Previous studies linking these concepts fail to include control conditions with other similar emotions (e.g., anger, fear), or use improper statistical techniques that inflate the appearance of specialness (Cameron, Lindquist, & Gray, in press). In our comprehensive review, we found only one study that met the threshold for experimental controls, and it found no evidence for a special disgust purity link (Cheng, Ottati, & Price, 2013)—consistent with an overarching dyadic template.

### *The Evolution of Moral Theory*

Our morals change over time. Smoking used to be a matter of personal preference, but now seems morally tinged. Pre-marriage cohabitation used to be deeply sinful, but now seems like a prudent choice before lifelong commitment. Likewise, the scientific understanding of morality changes over time. How people divide norm violations into “inappropriate” versus “immoral,” was once thought to be a matter of both harm and rational reason. Years later, it was discovered that moral judgment was driven by intuitions and emotion (Haidt, 2001). However, in their rush to abandon reason, these intuitive theories also largely abandoned harm. Inspired by classic research on categorization (Murphy, 2004), anthropological pluralism (Shweder, 2012),

and modern accounts of cognition (Cameron et al., in press), dyadic morality has reinstalled harm to its rightful place.

A flurry of recent research reveals that harm is both central to moral judgments (Schein & Gray, in press), and also shapes our perception of the moral world (Gray et al., 2014). This same research also reveals that moral differences between people do not reflect deep cognitive differences: Whether British or Indian, liberal or conservative, people have the same harm-based template based upon the perception of two minds—and it is these perceptions that drive moral debate. Whether people are pro-life or pro-choice, or for or against gay rights, depends upon the same simple—but ambiguous—question: do they perceive a mind being harmed?

The perception of harm may be subjective, but it has the hot truth of reality to those who see it. The “truthiness” of these perceptions suggests that arguing about perceived harm may be as difficult as arguing about the sexiness of fetishes. Depending on whom you ask, watching consenting adults spanking, screaming, and biting each other may be arousing or morally repugnant. As the British Board of Film Censors (BBFC) suggests, the issue of censorship does indeed depend on the question of whether children are harmed. But answers to the question of harm—like those of sexiness—are in the eye of the beholder.

## References

- Baron-Cohen, S. (2011). *The Science of Evil: On Empathy and the Origins of Cruelty*. New York, NY: Basic Books.
- Cameron, C. D., Lindquist, K. A., & Gray, K. (in press). A constructionist review of morality and emotions: No evidence for specific links between moral content and discrete emotions. *Personality and Social Psychology Review*.
- Cheng, J. S., Ottati, V. C., & Price, E. D. (2013). The arousal model of moral condemnation. *Journal of Experimental Social Psychology, 49*(6), 1012–1018. <http://doi.org/10.1016/j.jesp.2013.06.006>
- Gray, K., & Keeney, J. E. (2015). Impure, or just weird? Scenario sampling bias raises questions about the foundation of moral cognition. *Social Psychological and Personality Science*.
- Gray, K., Schein, C., & Ward, A. F. (2014). The myth of harmless wrongs in moral cognition: Automatic dyadic completion from sin to suffering. *Journal of Experimental Psychology: General, 143*(4), 1600–1615. <http://doi.org/10.1037/a0036149>
- Gray, K., Waytz, A., & Young, L. (2012). The moral dyad: A fundamental template unifying moral judgment. *Psychological Inquiry, 23*, 206–215.
- Greene, J. D. (2013). *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. Penguin.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*(4), 814–834.
- Haidt, J. (2012). *The righteous mind: Why good people are divided by politics and religion*. New York, NY: Pantheon Books.
- Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In T. Mischel (Ed.), *Cognitive development and epistemology* (pp. 151–235). New York, NY: Academic Press.
- Murphy, G. L. (2004). *The big book of concepts*. Cambridge, MA: MIT Press.
- Patterson, C. J. (2006). Children of Lesbian and Gay Parents. *Current Directions in Psychological Science, 15*(5), 241–244. <http://doi.org/10.1111/j.1467-8721.2006.00444.x>
- Quinn, B. (2014, December 2). Pornography law bans list of sexual acts from UK-made online films. Retrieved June 16, 2015, from <http://www.theguardian.com/culture/2014/dec/02/pornography-law-bans-list-sexual-acts-uk-made-online-films>

Schein, C., & Gray, K. (in press). The unifying moral dyad: Liberals and conservatives share the same harm-based moral template. *Personality & Social Psychology Bulletin*.

Shweder, R. A. (2012). Relativism and Universalism. In D. Fassin (Ed.), *A Companion to Moral Anthropology* (pp. 85–102). John Wiley & Sons, Ltd.

Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge, UK: Cambridge University Press.