

# Lecture 08 : Moral Psychology

Stephen A. Butterfill  
< s.butterfill@warwick.ac.uk >

Thursday, 12th March 2026

## Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Dual Process Theories: Objections, Evidence and Significance</b>          | <b>2</b>  |
| 1.1      | Background . . . . .   | 2         |
| 1.2      | Significance . . . . .   | 2         |
| <b>2</b> | <b>Evidence for Dual Process Theories</b>                                    | <b>3</b>  |
| 2.1      | Aim . . . . .  | 3         |
| 2.2      | Three Studies and Their Predictions . . . . .                                | 4         |
| 2.3      | What Did They Find and What Are Their Limits? . . . . .                      | 4         |
| 2.4      | Next Steps . . . . .   | 5         |
| <b>3</b> | <b>Conflicting Evidence against a Dual-Process Theory of Moral Judgement</b> | <b>5</b>  |
| 3.1      | Time Pressure . . . . .  | 6         |
| 3.2      | Process Dissociation . . . . .   | 7         |
| 3.3      | Conflicts in the Conflicting Evidence . . . . .                              | 7         |
| 3.4      | Conclusion . . . . .   | 8         |
| 3.5      | Appendix: Some Other Evidence . . . . .                                      | 8         |
| <b>4</b> | <b>Conclusion: No End to Our Troubles</b>                                    | <b>9</b>  |
| <b>5</b> | <b>Appendix: Dual Process Theory and Auxiliary Hypotheses</b>                | <b>10</b> |
| 5.1      | What Is the Core Idea . . . . .  | 11        |
| 5.2      | What Are the Candidate Auxiliary Hypotheses? . . . . .                       | 11        |
| 5.3      | Which Auxiliary Hypothesis to Choose? . . . . .                              | 12        |
|          | <b>Glossary</b>  | <b>12</b> |

# 1. Dual Process Theories: Objections, Evidence and Significance

A brief introduction to the three topics of this lecture.

This week's lecture builds on *Lecture 07*. Although you don't need to understand everything in that lecture, most of this week's lecture is about the loose reconstruction of Greene (2014)'s argument (see *Greene contra Ethics* in *Lecture 07*).

## 1.1. Background

The essence of many debunking arguments is captured by Railton (who is offering considerations with which to counter them):

‘a better understanding of the [...] origin of “intuitive” moral judgments might show them to be something other than manifestations of underlying moral competencies or principles. [...] “moral intuitions” might therefore deserve less deference [...] than they characteristically receive in philosophical [...] moral thought’ (Railton 2014, p. 832).

~832]{railton:2014\_affective}. \end{quote}

The argument we are evaluating (see *Greene contra Ethics* in *Lecture 07*) concludes that:

Not-justified-inferentially premises about particular moral scenarios cannot be used in ethical arguments where the aim is to establish knowledge of their conclusions.

This argument differs from Railton's sketch in that it is based on considerations which support, rather than undermine, the view that moral intuitions are manifestations of an underlying competence. That competence is, however, limited in ways which mean we should not always rely on it in philosophical arguments (*Cognitive Miracles: When Are Fast Processes Unreliable?* in *Lecture 07*).

## 1.2. Significance

Why is the debunking argument we are considering (see *Greene contra Ethics* in *Lecture 07*) significant?

One reason is that it creates complications for views about the value of moral intuitions in doing ethics. Consider Audi's view:

‘Intuition is a resource in all of philosophy, but perhaps nowhere more than in ethics’ (Audi 2015, p. 57).

‘Episodic intuitions [...] can serve as data [...] ... beliefs that derive from them receive prima facie justification.’ (Audi 2015, p. 65).

Audi also compares episodic intuitions to perceptions:

‘episodic intuitions evidence propositions in a non-inferential way that [...] resembles] the way in which perceptions evidence propositions’ (Audi 2015, p. 74).

The case of vertical motion illustrates that perceptions do not actually justify beliefs about unfamiliar situations, such as those involving spiral tubes or vertical motion (as we saw in *Linking Ethics to Moral Psychology: Dual-Process Theories* in Lecture 07). Similarly, if the argument we are evaluating succeeds, and if Audi’s episodic intuitions are consequences of faster processes, then your knowledge of the limits of faster processes defeats any prima facie justification which moral intuitions might provide for beliefs about unfamiliar situations. Whatever your intuitions, they do not enable you to know that it is wrong to sacrifice one person to save five in Footbridge.

## 2. Evidence for Dual Process Theories

What is the strongest evidence in favour of our stripped-down dual-process theory of moral psychology (see *A Dual Process Theory of Ethical Judgement* in Lecture 07)? Greene (2014) cites many studies. In this section we evaluate three of them, including one involving process dissociation (Conway & Gawronski 2013). (As a bonus, process dissociation also enables us to revisit the issue of whether emotion influences moral judgement.)

### 2.1. Aim

The first premise of our stripped-down dual process theory of moral psychology (see *A Dual Process Theory of Ethical Judgement* in Lecture 07) is probably the most controversial:

1. Ethical judgements are explained by a dual-process theory, which distinguishes faster from slower processes.

What is the strongest evidence in favour of this premise? Greene (2014) cites many studies. As always we should not take for granted that Greene’s description of the studies is correct: we need to evaluate them for ourselves

(see the step-by-step guide in *Moral Intuitions and Emotions: Evaluating the Evidence* in Lecture 02).

We are looking for evidence in favour of the stripped-down dual process theory together with our selected auxiliary hypothesis:

Auxiliary Assumption: Only the slow process flexibly takes into account differences in the more distal outcomes of an action (see *Appendix: Dual Process Theory and Auxiliary Hypotheses* (section §5)).

## 2.2. Three Studies and Their Predictions

Here we will consider three of the studies Greene cites. It is important to specify which prediction each study tests (which may not be obvious from the abstract).

- Suter & Hertwig (2011) – prediction: limiting the time available to make a decision will reduce the influence of distal outcomes.
- Trémolière & Bonnefon (2014) – prediction: limiting the time available to make a decision will reduce consequentialist responses.<sup>1</sup>
- Conway & Gawronski (2013) – prediction: higher cognitive load will reduce the dominance of the more outcome-sensitive process.

## 2.3. What Did They Find and What Are Their Limits?

Suter & Hertwig (2011) is an example of a relatively simple study which provides evidence in favour of the dual process theory plus auxiliary hypothesis.

One limit of this study is that it does not involve any variation in the distal outcomes of actions. This is relevant because the auxiliary hypothesis is about how different processes are differently influenced by distal outcomes.

Although not designed with exactly this in mind, Trémolière & Bonnefon (2014) does observe responses to otherwise similar actions with different distal outcomes. However, the findings are not predicted by the dual process

---

<sup>1</sup> As explained in the lecture recording, this study is associated with a second prediction, which the results appear to disconfirm: limiting the time available to make a decision will reduce sensitivity to outcomes.

theory and auxiliary hypothesis.<sup>2</sup>

One limit of both Suter & Hertwig (2011) and Trémolière & Bonnefon (2014) is that they treat responses as either consequentialist or not. These studies are sometimes presented as comparing consequentialist with deontological responses; but this cannot be accurate because failing to respond as a consequentialist does not make you a deontologist (you may be neither).

Conway & Gawronski (2013) overcome this limit in addition to observing responses to otherwise similar actions with different distal outcomes. It is one of the strongest tests of the stripped-down dual process theory and its auxiliary hypothesis. These authors find, as predicted, that higher cognitive load reduces sensitivity to outcomes while not affecting sensitivity to moral prohibitions (such as on killing).

Conway & Gawronski (2013) are also important because they introduce process dissociation in moral psychology. Although difficult to understand (I attempt to explain the bare minimum you need in the lecture), this is a powerful method for testing theories.

## 2.4. Next Steps

Having evaluated some of the evidence in favour of the dual process theory, our next task is to consider evidence against it.

## 3. Conflicting Evidence against a Dual-Process Theory of Moral Judgement

Several studies provide results which partially undermine the evidence in favour of our stripped-down dual-process theory of moral cognition. Here we consider two of the most compelling (Bago & De Neys 2019; Gawronski et al. 2017). Taken together these studies are puzzling: as well as individually conflicting with the evidence for our dual-process theory, the two studies also appear to conflict with each other. It is hard to identify a view that is consistent with taking the results from all of the studies at face value.

We have a stripped down dual-process theory of moral judgement (see *A Dual Process Theory of Ethical Judgement* in Lecture 07) and an auxiliary hypothesis (see *Appendix: Dual Process Theory and Auxiliary Hypotheses* (section §5)). According to these:

---

<sup>2</sup> Although the Trémolière & Bonnefon (2014)'s findings may be interpreted as disconfirming a prediction (as Gawronski & Beer 2017, p. 669 propose), it would be incautious to rely on post hoc reinterpretations of findings.

Two (or more) ethical processes are distinct in this sense: the conditions which influence whether they occur, and which outputs they generate, do not completely overlap.

One process is faster than another: it makes fewer demands on scarce cognitive resources such as attention, inhibitory control and working memory.

Only the slow process flexibly takes into account differences in the more distal outcomes of an action.

Earlier we saw that there is some evidence which appears to support the predictions of this theory (in *Evidence for Dual Process Theories* (section §2)). Is there also evidence disconfirming any of its predictions?

While it is hard to find evidence directly against this theory,<sup>3</sup> there are some studies that undermine the view we took earlier on which studies provide evidence in favour of the dual-process theory.

### 3.1. Time Pressure

Recall that Suter & Hertwig (2011) provide evidence that time pressure makes participants less sensitive to distal outcomes. Bago & De Neys (2019) consider what happens when subjects first make a moral judgement under time pressure and extraneous cognitive load and then, just after, make another moral judgement (in answer to the same question) with no time pressure and no extraneous cognitive load. They report:

‘Our critical finding is that although there were some instances in which deliberate correction occurred, these were the exception rather than the rule. Across the studies, results consistently showed that in the vast majority of cases in which people opt for a [consequentialist] response after deliberation, the [consequentialist] response is already given in the initial phase’ (Bago & De Neys 2019, p. 1794).

Bago & De Neys (2019)’s findings suggest that time pressure and cognitive load do not make people less consequentialist, nor more consequentialist. This is an obstacle to considering Suter & Hertwig (2011)’s study as evidence for our dual-process theory of moral judgement.

Further studies make the picture even more complex. Rosas & Aguilar-Pardo (2020) observed the converse to Suter & Hertwig (2011): in their studies, lim-

---

<sup>3</sup> One potential source of evidence that directly opposes the theory is Bialek & De Neys (2017) (mentioned below). Unfortunately I came across this too late to include it in the recording.

iting the time available to make a decision increases characteristically consequentialist responses.<sup>4</sup> Vega et al. (2021) found faster responses linked to deontological responses in their Study 1 and the opposite pattern in their Study 2, where both studies use the same paradigm (one of the goals of Study 2 was to replicate the results of Study 1; p. 131).

### 3.2. Process Dissociation

Recall that Conway & Gawronski (2013) use process dissociation to provide evidence for the prediction that higher cognitive load reduces sensitivity to more distal outcomes.

Gawronski et al. (2017) note that reduced sensitivity to more distal outcomes could be consequence of a general preference not to do anything when under time pressure. They therefore extend the process dissociation model to include a preference for no action.

Separating sensitivity to distal outcomes from preferences not to act changes the picture:

‘The only significant effect in these studies was a significant increase in participants’ general preference for inaction as a result of cognitive load. Cognitive load did not affect participants’ sensitivity to morally relevant consequences’ (Gawronski et al. 2017, p. 363).

They conclude:

‘cognitive load influences moral dilemma judgments by enhancing the omission bias, not by reducing sensitivity to consequences in a utilitarian sense’ (Gawronski et al. 2017, p. 363).

While we should be cautious about putting too much weight on this study, these results do reveal that we cannot take Conway & Gawronski (2013) as evidence in favour of our dual-process theory and auxiliary hypothesis.<sup>5</sup>

### 3.3. Conflicts in the Conflicting Evidence

The two studies which conflict with the evidence for our dual-process theory also appear to conflict with each other. If Gawronski et al. (2017) is

---

<sup>4</sup> Kurzban et al. (2012, p. 331) provide theoretical justification for the possibility that faster processes are characteristically consequentialist and it is only the slower processes that enable humans to act in accordance with deontological principles.

<sup>5</sup> As an aside for fans of process dissociation, Qian et al. (2024) built on Gawronski et al. (2017)’s work and found some unexpected and interesting differences between gender and nationality (Japanese, Chinese and US respondents).

right about cognitive load, the participants in Bago & De Neys (2019)'s study should have appeared to be less 'utilitarian' (as they describe it) when under cognitive load. This is because avoiding action would lead one to make judgements that Bago & De Neys classify as non-utilitarian.

So we cannot accept both Gawronski et al. (2017)'s and Bago & De Neys (2019)'s conclusions.

This is a sign that there may be something wrong with the way the studies are constructed, perhaps because the dual-process theories they are targeting are not well specified (e.g. involve too many independent bets being made simultaneously).

### 3.4. Conclusion

We may not yet have found sufficient grounds to reject the stripped-down dual-process theory of moral cognition outright. But we should recognise that we do not have sufficient evidence to confidently assert that any of the candidate auxiliary hypotheses are true (see *Appendix: Dual Process Theory and Auxiliary Hypotheses* (section §5)). On balance, it seems reasonable to deviate from the mainstream in not accepting any auxiliary hypothesis as firmly established while provisionally accepting, in line with the mainstream, that the stripped-down dual-process theory (see *A Dual Process Theory of Ethical Judgement* in Lecture 07) will turn out to be true.

This matters for Greene (2014)'s attempt to link characteristically consequentialist judgements to slow processes. As things stand, we do not know that any such link exists. We should be correspondingly cautious in using the dual-process theory in defending a consequentialist ethical theory.

### 3.5. Appendix: Some Other Evidence

There is much evidence on how time pressure and cognitive load influence moral judgements. Understanding how it bears on the stripped-down dual process theory is complicated, in part because many studies target features of Greene's dual process theory that are not features of the stripped-down dual process theory. Here my focus is on studies that can be interpreted as finding evidence against the theory.

Bialek & De Neys (2017) provide direct evidence against our auxiliary hypothesis: time pressure and cognitive load do not appear to influence the extent to which participants take into account the distal outcomes of an action in making moral judgements.

Tinghög et al. (2016) find no evidence for effects of time pressure or cognitive load on moral judgements. They conclude that 'intuitive moral decision-

making does not differ from decisions made in situations where deliberation before decision is facilitated.’

Baron & Gürçay (2017) offer a meta-analysis of response time findings, but focus on an auxiliary hypothesis which we have not used (the ‘default interventionist’ claim).

Koop (2013) and Gürçay & Baron (2017) both measure subjects’ movements as they make a decision, which can provide a window on to how the decision unfolds. Koop (2013) do not find evidence to support the conjecture that subjects increasingly consider distal outcomes later in the decision process. Gürçay & Baron (2017) do not find support for the conjecture that more thinking increases sensitivity to the distal outcomes of actions.

Capraro et al. (2019) examined the effects of telling (they say ‘priming’) people to use ‘emotion, rather than reason’. As background, they note that much of the research on dual-process theories concerns characteristically consequentialist judgements, which may confound two factors: reluctance to cause harm instrumentally and impartiality. The auxiliary hypothesis we have chosen is linked to the first factor (reluctance to cause harm instrumentally) but not the second. They find that when these factors are separated, priming intuition reduces willingness to cause harm instrumentally.<sup>6</sup>

Although Capraro et al. (2019)’s study supports the auxiliary hypothesis, I have included it here (in a section on evidence against our dual-process theory of moral judgement) because it illustrates a complication in interpreting studies which appear to provide evidence against the theory: none of them are focussed narrowly on sensitivity to distal outcomes specifically rather than on some broader contrast between characteristically consequentialist and characteristically deontological.

Rehren (2024) offers a meta-analysis of studies which manipulate cognitive load and time pressure (among other factors). While I am unsure whether we should be convinced by the findings of the meta-analysis, Table 1 and Figure 2 are helpful as an index of a wide range of studies.

## 4. Conclusion: No End to Our Troubles

The dual process theory of ethical judgement appears, so far, to be well supported by evidence for it (although we have yet to consider any evidence

---

<sup>6</sup> Bartels (2008) distinguished between subjects with more intuitive and more deliberative thinking styles. He found that moral judgement ‘(a) makes use of intuitive and deliberative process, (b) is influenced by the judgment-eliciting context, and (c) recruits representations of both deontological constraints and utilitarian considerations.’

against it). This supports the loose reconstruction of Greene (2014)'s argument (see *Greene contra Ethics* in Lecture 07), which is also not vulnerable to any of the quick objections (at least none I could find in the literature). We may therefore have to accept the conclusion that not-justified-inferentially premises about particular moral scenarios, and about debatable ethical principles, cannot be used in ethical arguments insofar as the arguments aim to establish knowledge of their conclusions. This does not show that consequentialism is the one true ethical theory. But it does imply that we should avoid a variety of approaches to doing ethics, including Foot's, Kamm's and Rawls' method of reflective equilibrium insofar as our aim is to gain knowledge of ethical truths.

## 5. Appendix: Dual Process Theory and Auxiliary Hypotheses

The stripped-down dual process theory (see *A Dual Process Theory of Ethical Judgement* in Lecture 07) requires at least one auxiliary hypothesis in order to relate it to available evidence. At least six auxiliary hypotheses have been proposed. Which should we accept?

*This is not covered in the lecture this year. You do not need this material, but you may wish to consider it if you are writing about dual-process theories of ethical judgement.*

Proponents of dual process theories of ethical judgments tend to offer complex sets of hypotheses which go far beyond the core idea of a dual process theory and are not all equally well supported by evidence (for two examples, see Greene and Kumar in *A Dual Process Theory of Ethical Judgement* in Lecture 07).

Since the hypotheses that are associated with different proposals can be isolated and tested separately, combining them yields something more like a parlay bet than a theory.<sup>7</sup> Better to keep them separate.

---

<sup>7</sup> This is too quick: all theories are bets. And I don't have anything against parlay bets—if you are given good enough odds, they can be spectacular (and fun, of course). But making progress in understanding is like gambling in that it requires managing risk. Just as you have to allocate cash to bets carefully in order to minimize the risk of ruin, so in your research you have to allocate time carefully in order to avoid being left with nothing. (There is some research on how this is done that starts with the Kelly Criterion.) You should not put everything into a small number of especially risky bets. This is the true problem with Parlay bets.

## 5.1. What Is the Core Idea

Just this and nothing more.

Two (or more) ethical processes are distinct in this sense: the conditions which influence whether they occur, and which outputs they generate, do not completely overlap.

One process is faster than another: it makes fewer demands on scarce cognitive resources such as attention, inhibitory control and working memory.

Because much of the available evidence involves observations of just one response to a moral scenario (typically a verbal judgement<sup>8</sup>), additional, auxiliary hypotheses are needed to generate predictions linked to the currently available evidence. But which auxiliary hypothesis should we adopt?

## 5.2. What Are the Candidate Auxiliary Hypotheses?

At least six candidates for auxiliary hypotheses can be found in the existing literature (and there are probably many more):

| id | fast                                | slow                         | source                       |
|----|-------------------------------------|------------------------------|------------------------------|
| 1. | never consequentialist <sup>9</sup> | sometimes consequentialist   | Greene (2014)                |
| 2. | always deontological <sup>10</sup>  | sometimes not deontological  | Greene (2014) <sup>11</sup>  |
| 3. | always only act-types               | sometimes not only act-types | Cushman (2013)               |
| 4. | never distal outcomes               | sometimes distal outcomes    | Cushman (2013) <sup>12</sup> |

<sup>8</sup> Measures like response time, movement trajectory, proactive eye gaze or pupil dilation do not seem to have caught on in moral psychology despite being used to good effect in supporting dual process theories in other domains including social cognition (see van der Wel et al. 2014 or Edwards & Low 2017 for some examples).

<sup>9</sup> In this context, the term ‘consequentialist’ is short for characteristically consequentialist.

<sup>10</sup> In this context, the term ‘deontological’ is short for characteristically deontological.

<sup>11</sup> Greene (2014) does not explicitly separate this and the first candidate, (1). But they are clearly distinct.

<sup>12</sup> Cushman (2013) phrases the distinction in terms of actions and outcomes. This is potentially confusing if, like Davidson (1971), you think that actions are individuated by their outcomes. But it is clear from the discussion that Cushman’s distinction concerns a contrast between more proximal outcomes and more distal outcomes. Note also that Cushman does not explicitly distinguish this and the third candidate, (3); perhaps the distinction is too minor to consider.

| id | fast       | slow        | source         |
|----|------------|-------------|----------------|
| 5. | model-free | model-based | Cushman (2013) |
| 6. | affective  | cognitive   | Greene (2014)  |

**UPDATE** A further candidate auxiliary hypothesis might concern the distinction between personal and impersonal. Since this introduces even more problems than we will eventually have to face (McGuire et al. 2009), it does not appear to be a good bet as things stand.

### 5.3. Which Auxiliary Hypothesis to Choose?

I propose that we use just one auxiliary hypothesis, which is a version of (4) above:

Only the slow process flexibly takes into account the more distal outcomes of an action.

This auxiliary hypothesis has two virtues. Together with the stripped-down dual process theory, it generates nearly all the predictions for which we currently have good evidence. And it cannot be significantly weakened while still doing this.

## Glossary

**automatic** As we use the term, a process is *automatic* just if whether or not it occurs is to a significant extent independent of your current task, motivations and intentions. To say that *mindreading is automatic* is to say that it involves only automatic processes. The term ‘automatic’ has been used in a variety of ways by other authors: see Moors (2014, p. 22) for a one-page overview, Moors & De Houwer (2006) for a detailed theoretical review, or Bargh (1992) for a classic and very readable introduction 13

**characteristically consequentialist** According to Greene, a judgement is *characteristically consequentialist* (or ‘characteristically utilitarian’) if it is one in ‘favor of characteristically consequentialist conclusions (eg, “Better to save more lives”)’ (Greene 2007, p. 39). According to Gawronski et al. (2017, p. 365), ‘a given judgment cannot be categorized as [consequentialist] without confirming its property of being sensitive to consequences.’ 8, 9, 11

**characteristically deontological** According to Greene, a judgement is *characteristically deontological* if it is one in ‘favor of characteristically deontological conclusions (eg, “It’s wrong despite the benefits”)’ (Greene 2007, p. 39). According to Gawronski et al. (2017, p. 365), ‘a given judgment cannot be categorized as deontological without confirming its property of being sensitive to moral norms.’ 9, 11

**cognitively efficient** A process is *cognitively efficient* to the degree that it does not consume working memory, attention, inhibitory control or other scarce cognitive resources. 13

**debunking argument** A *debunking argument* aims to use facts about why people make a certain judgement together with facts about which factors are morally relevant in order to undermine the case for accepting it. Königs (2020, p. 2607) provides a useful outline of the logic of these arguments (which he calls ‘arguments from moral irrelevance’): ‘when we have different intuitions about similar moral cases, we take this to indicate that there is a moral difference between these cases. This is because we take our intuitions to have responded to a morally relevant difference. But if it turns out that our case-specific intuitions are responding to a factor that lacks moral significance, we no longer have reason to trust our case-specific intuitions suggesting that there really is a moral difference. This is the basic logic behind arguments from moral irrelevance’ (Königs 2020, p. 2607). 2

**distal outcome** The outcomes of an action can be partially ordered by the cause-effect relation. For one outcome to be more *distal* than another is for it to be lower with respect to that partial ordering. To illustrate, if you kick a ball through a window, the window’s breaking is a more distal outcome than the kicking. 4, 6–9, 11, 12

**dual-process theory** Any theory concerning abilities in a particular domain on which those abilities involve two or more processes which are distinct in this sense: the conditions which influence whether one process occurs differ from the conditions which influence whether another occurs. 3, 5, 8

**fast** A *fast* process is one that is to some interesting degree cognitively efficient (and therefore likely also some interesting degree automatic). These processes are also sometimes characterised as able to yield rapid responses.

Since automaticity and cognitive efficiency are matters of degree, it is only strictly correct to identify some processes as faster than others.

The fast–slow distinction has been variously characterised in ways that do not entirely overlap (even individual authors have offered differing characterisations at different times; e.g. Kahneman 2013; Morewedge & Kahneman 2010; Kahneman & Klein 2009; Kahneman 2002): as its advocates stress, it is a rough-and-ready tool rather than an element in a rigorous theory. 3, 6, 11, 15

**Footbridge** A dilemma; also known as *Drop*. A runaway trolley is about to run over and kill five people. You can hit a switch that will release the bottom of a footbridge and one person will fall onto the track. The trolley will hit this person, slow down, and not hit the five people further down the track. Is it okay to hit the switch? 3

**loose reconstruction** (of an argument). A reconstruction which prioritises finding a correct argument for a significant conclusion over faithfully representing the argument being reconstructed. 2

**moral intuition** According to this lecturer, a person's intuitions are the claims they take to be true independently of whether those claims are justified inferentially. And a person's *moral* intuitions are simply those of their intuitions that concern ethical matters.

According to Sinnott-Armstrong et al. (2010, p. 256), moral intuitions are 'strong, stable, immediate moral beliefs.'

Audi (2015) distinguishes various notions of intuition including episodic intuitions. These 'episodic intuitions' are supposed to be analogous to perceivings: they are 'intellectual seemings ... of the truth of a proposition.' 2, 3

**not-justified-inferentially** A claim (or premise, or principle) is not-justified-inferentially if it is not justified in virtue of being inferred from some other claim (or premise, or principle).

Claims made on the basis of perception ('That jumper is red', say) are typically not-justified-inferentially.

Why not just say 'noninferentially justified'? Because that can be read as implying that the claim *is* justified, noninferentially. Whereas 'not-justified-inferentially' does not imply this. Any claim which is not justified at all is thereby not-justified-inferentially. 2, 10

**outcome** An outcome of an action is a possible or actual state of affairs. 13, 15

**proximal outcome** The outcomes of an action can be partially ordered by the cause-effect relation. For one outcome to be more *proximal* than another is for it to be higher with respect to that partial ordering. To illustrate, if you kick a ball through a window, the kicking is a more proximal outcome than the window's breaking. 11

slow converse of fast. 3

## References

- Audi, R. (2015). Intuition and Its Place in Ethics. *Journal of the American Philosophical Association*, 1(1), 57–77.
- Bago, B. & De Neys, W. (2019). The Intuitive Greater Good: Testing the Corrective Dual Process Model of Moral Cognition. *Journal of Experimental Psychology: General*, 148(10), 1782–1801.
- Bargh, J. A. (1992). The Ecology of Automaticity: Toward Establishing the Conditions Needed to Produce Automatic Processing Effects. *The American Journal of Psychology*, 105(2), 181–199.
- Baron, J. & Gürçay, B. (2017). A meta-analysis of response-time tests of the sequential two-systems model of moral judgment. *Memory & Cognition*, 45(4), 566–575.
- Bartels, D. M. (2008). Principled moral sentiment and the flexibility of moral judgment and decision making. *Cognition*, 108(2), 381–417.
- Bialek, M. & De Neys, W. (2017). Dual processes and moral conflict: Evidence for deontological reasoners' intuitive utilitarian sensitivity. *Judgment and Decision Making*, 12(2), 148.
- Capraro, V., Everett, J. A. C., & Earp, B. D. (2019). Priming intuition disfavors instrumental harm but not impartial beneficence. *Journal of Experimental Social Psychology*, 83, 142–149.
- Conway, P. & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of personality and social psychology*, 104(2), 216–235.
- Cushman, F. (2013). Action, Outcome, and Value: A Dual-System Framework for Morality. *Personality and Social Psychology Review*, 17(3), 273–292.
- Davidson, D. (1971). Agency. In R. Binkley, R. Bronaugh, & A. Marras (Eds.), *Agent, Action, and Reason*, (pp. 3–25). Toronto: University of Toronto Press.

- Reprinted in Davidson, D. (1980) *Essays on Actions and Events*. Oxford: Oxford University Press.
- Edwards, K. & Low, J. (2017). Reaction time profiles of adults' action prediction reveal two mindreading systems. *Cognition*, *160*, 1–16.
- Gawronski, B., Armstrong, J., Conway, P., Friesdorf, R., & Hütter, M. (2017). Consequences, norms, and generalized inaction in moral dilemmas: The CNI model of moral decision-making. *Journal of personality and social psychology*, *113*(3), 343–376.
- Gawronski, B. & Beer, J. S. (2017). What makes moral dilemma judgments “utilitarian” or “deontological”? *Social Neuroscience*, *12*(6), 626–632.
- Greene, J. D. (2007). The Secret Joke of Kant's Soul. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Vol. 3* (pp. 35–79). MIT Press.
- Greene, J. D. (2014). Beyond Point-and-Shoot Morality: Why Cognitive (Neuro)Science Matters for Ethics. *Ethics*, *124*(4), 695–726.
- Gürçay, B. & Baron, J. (2017). Challenges for the sequential two-system model of moral judgement. *Thinking & Reasoning*, *23*(1), 49–80.
- Kahneman, D. (2002). Maps of bounded rationality: A perspective on intuitive judgment and choice. In T. Frangsmyr (Ed.), *Le Prix Nobel, ed. T. Frangsmyr, 416–499.*, volume 8 (pp. 351–401). Stockholm, Sweden: Nobel Foundation.
- Kahneman, D. (2013). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kahneman, D. & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, *64*(6), 515–526.
- Königs, P. (2020). Experimental ethics, intuitions, and morally irrelevant factors. *Philosophical Studies*, *177*(9), 2605–2623.
- Koop, G. J. (2013). An assessment of the temporal dynamics of moral decisions. *Judgment and decision making*, *8*(5), 527.
- Kurzban, R., DeScioli, P., & Fein, D. (2012). Hamilton vs. Kant: Pitting adaptations for altruism against adaptations for moral judgment. *Evolution and Human Behavior*, *33*(4), 323–333.
- McGuire, J., Langdon, R., Coltheart, M., & Mackenzie, C. (2009). A reanalysis of the personal/impersonal distinction in moral psychology research. *Journal of Experimental Social Psychology*, *45*(3), 577–580.

- Moors, A. (2014). Examining the mapping problem in dual process models. In *Dual process theories of the social mind* (pp. 20–34). Guilford.
- Moors, A. & De Houwer, J. (2006). Automaticity: A Theoretical and Conceptual Analysis. *Psychological Bulletin*, *132*(2), 297–326.
- Morewedge, C. K. & Kahneman, D. (2010). Associative processes in intuitive judgment. *Trends in Cognitive Sciences*, *14*(10), 435–440.
- Qian, Y., Takimoto, Y., Wang, L., & Yasumura, A. (2024). Exploring cultural and gender differences in moral judgment: A cross-cultural study based on the CNI model. *Current Psychology*, *43*(6), 5243–5253. [Online; accessed 2025-03-11].
- Railton, P. (2014). The Affective Dog and Its Rational Tale: Intuition and Attunement. *Ethics*, *124*(4), 813–859.
- Rehren, P. (2024). The effect of cognitive load, ego depletion, induction and time restriction on moral judgments about sacrificial dilemmas: a meta-analysis. *Frontiers in Psychology*, *15*. [Online; accessed 2025-03-04].
- Rosas, A. & Aguilar-Pardo, D. (2020). Extreme time-pressure reveals utilitarian intuitions in sacrificial dilemmas. *Thinking & Reasoning*, *26*(4), 534–551.
- Sinnott-Armstrong, W., Young, L., & Cushman, F. (2010). Moral intuitions. In J. M. Doris, M. P. R. Group, et al. (Eds.), *The moral psychology handbook* (pp. 246–272). Oxford: Oxford University Press.
- Suter, R. S. & Hertwig, R. (2011). Time and moral judgment. *Cognition*, *119*(3), 454–458.
- Tinghög, G., Andersson, D., Bonn, C., Johannesson, M., Kirchler, M., Koppel, L., & Västfjäll, D. (2016). Intuition and Moral Decision-Making – The Effect of Time Pressure and Cognitive Load on Moral Judgment and Altruistic Behavior. *PLOS ONE*, *11*(10), e0164012.
- Trémolière, B. & Bonnefon, J.-F. (2014). Efficient Kill–Save Ratios Ease Up the Cognitive Demands on Counterintuitive Moral Utilitarianism. *Personality and Social Psychology Bulletin*, *124*(3), 379–384.
- van der Wel, R. P. R. D., Sebanz, N., & Knoblich, G. (2014). Do people automatically track others’ beliefs? evidence from a continuous measure. *Cognition*, *130*(1), 128–133.
- Vega, S., Mata, A., Ferreira, M. B., & Vaz, A. R. (2021). Metacognition in moral decisions: judgment extremity and feeling of rightness in moral intuitions. *Thinking & Reasoning*, *27*(1), 124–141.