The Fate of the Third Chimpanzee

The Nicod Lectures 2008

Kim Sterelny
Philosophy and Tempo and Mode
Victoria University of Wellington and the ANU

Session 2

Moral Nativism: A Sceptical Response

I

Introduction

In the first session, I noted that the standard model takes language to be a general template for cognitive competence. We reliably acquire our abilities to navigate in our social, psychological, normative and technological worlds only through having critical principles pre-wired in our minds, though often these are schematic or only partially specified. In their recent work on moral cognition, John Mikhail, Susan Dywer and Marc Hauser explicitly base their theory of moral cognition on this linguistic template. These ideas are the focus of this session. As it develops, I have three aims. One is critical. I do not think the linguistic template is a helpful way of thinking about moral cognition. A second aim is to apply the apprentice learning template sketched in session 1 to moral cognition, thereby illustrating the power of informationally engineering learning environments. Linguistic nativism is often organised around the idea that the environment of language learning is informationally impoverished, and the moral grammarians have followed this lead. I could not be more sceptical. A parental generation (I shall argue) engineers the developmental environment of the generation of their children, thus making the cross-generation flow of information about norms and values much more reliable than it would otherwise be. A third aim is

to explore evolutionary models of culture, and of the interactions between biological and cultural inheritance. I argue (in contrast to some meme-sceptics) that human cultural worlds do support high fidelity cultural learning. Moreover, it supports high fidelity cultural learning of information relevant to high stakes decision-making. So I see cultural learning as a high fidelity, high volume inheritance system. But while high fidelity cultural learning is necessary for the cumulative evolution of complex adaptations, it is not sufficient. Moral cognition is a good vehicle for exploring both the power of cultural evolution to build adaptation, and the limits on that power. I begin by sketching the prima facie case for the linguistic image of moral cognition.

II

The Adapted Moraliser?

We are distinctively talking, technological and co-operative apes. But we are also moralising apes. No other living primate moralises; arguably, all normal adults do. Some cultures may not cleanly discriminate between norms of morality, religion and disgust. But as far as we know, humans in all cultures think normatively, not just descriptively, though. As normative cognition is both pervasive and pan-cultural, it is no surprise that adaptationist approaches to human cognition have been

extended to moral cognition. On the view developed in this work, moral cognition is a distinctive cognitive adaptation, for its evolution gave our ancestors greater access to the benefits of co-operation. Moral cognition adapted us to life in a complex social world in which managing co-operation was of central importance. More particularly, in different ways, Robert Frank, Marc Hauser and Richard Joyce have all suggested that in a world in which securing stable co-operation was increasingly important, our ancestral cognitive phenotypes would have left us with a weakness-of-will problem. Our primate minds would not have been able to resist the temptation to secure short-term benefit at the greater, but delayed, expense of fracturing trust. Moral judgement increases the salience of the prosocial choice, making it easier for us to choose wisely.

Moralising, then, is individually adaptive. But the capacity to moralise depended on the evolution of a dedicated cognitive specialisation akin to language. Language is independent of central, conscious cognitive processing. An agent does not have to decide to hear speech as language; it is automatic, mandatory. Agents parse sentences of their native language, recovering their organization, but they have no introspective access to information used by the mechanisms which take speech as input and deliver to conscious awareness an interpretation of what has been said. This information is tacit. It is not portable; it is not

available to drive other cognitive processes, and hence cannot be expressed as assertions or beliefs. These phenomena are explained by the fact that the normal development of the human mind results in a language module coming on stream.

Hauser, Dwyer and Mikhail suggest that the human capacity for moral judgment has a structure parallel to that of language. This suggestion is made plausible by the phenomenological similarities between moral cognition and linguistic cognition. Moral assessment is fast and automatic. We do not have to decide whether to evaluate a situation normatively when we read of the abduction of a three-year old girl. Moreover, normative thought is not just a universal feature of human society and of normal agency; there is reason to suspect that there are family resemblance relations between the moral systems of different cultures. As with language, the capacity for moral cognition seems to develop impressively early and with impressive robustness. Those who are sceptical of the moral grammarians' picture will need to explain why moral cognition seems modular, even though it is not.

Moreover, the moral nativists argue that moral cognition is not just automatic and universal; it is universal in surprising and subtle ways. Thus Hauser argues that there is striking evidence of cross-cultural uniformity in moral judgements. Most agents

draw an important moral distinction between acts and omissions; and they draw an important moral distinction between the foreseen but unintended consequences of actions, and those consequences that are both foreseen and intended. So in pursuit of the best overall outcome (for example: saving the most lives possible in the face of immanent diaster) it is morally acceptable to tolerate bad consequences that are foreseen but unintended. But it is not morally acceptable to intentionally do evil to block a still greater evil. Thus they judge in accordance with the so-called principle of double effect. So, for example, subjects are presented with scenarios in which an agent is faced between saving five lives at the cost of one. The cases differ in the kind of interventions which are necessary in order to save the five. In some interventions, the agent must directly kill the one (using him as a brake, buffer or obstruction). His death is the direct target of intervention. To save the five from the runaway tram, an innocent bystander must be pushed onto the tracks to slow the tram, giving the five time to escape. In others, the one dies, but as a side-effect. The efficacy of the intervention to save the five does not in itself depend on its being fatal to the one. In one scenario, the agent has access to a switching mechanism that can divert the trolley bus from one line to another. Unfortunately, there is an innocent bystander on the spare line too, but while his death is foreseen it is not intended, and so diverting the train is morally permissible.

In developing this grammatical model of morality, Hauser emphasises the subtlety and the apparent cultural invariance of the general principles that seem to underlie particular judgements. But he also makes much of the introspective opacity of these principles. While many agents judge in accordance with the principle of double affect, very few can consciously articulate it. Agents reliably make moral discriminations, regularly and predictably judging that some actions are morally appropriate and others are not. But they are typically unable to articulate the principles that guide these discriminations. Thus many of our operative moral principles are tacit, and this leads to a puzzle about moral development. Moral development confronts a poverty-of-the-stimulus problem. We must be pre-wired for specific moral principles, for we come to rely on those principles in our fast and unreflective moral reasoning. Yet they have never been articulated and taught in the period in which moral competence develops. The moral nativists accept that there must be a significant cultural input to moral development, but they also argue that there is a very significant and specifically moral innate component.

However, though there are important phenomenological similarities between moral cognition and language, there are as

well important phenomenological and functional differences between moral competence and paradigmatic modular competences. Vision, for example, is modular because vision is an online cognitive task. We use vision to guide our motion through a cluttered world. Much information that comes to us through vision has a short-shelf life. If it is relevant at all, it is relevant now. So the system has to be fast. It can be encapsulated because the visual discriminations we make typically depend on very general and very stable features of our physical environment. Likewise, if we are to understand others, we must understand them as they talk. So the adaptive rationale of perceptual modules depends on two features: first: the tasks which they support are urgent, and second: the features of environments on which their reliability depends are stable. Moral decision-making typically does not have this urgency: we very rarely have to judge and act within a fraction of a second. Our moral evaluations are typically aspects of offline cognition than online cognition. We evaluate scenarios in conversation and imagination. We imaginatively project ourselves into possible futures and pasts, in "mental time travel". So we do not just evaluate actions and circumstances we perceive, and when we do, we rarely have to act in milliseconds. Moral cognition is not urgent cognition.

I think there is an even more telling contrast, and that is the focus on the next section. Explicit moral beliefs play a central role in moral cognition. In contrast, while we have a some scattered folk metalinguistic views, they are not central to linguistic competence; to executing one's own linguistic intentions, or to understanding those of others. Many speakers have essentially no explicit beliefs about syntax or phonology at all. Moral thought is very different: we do not just have explicit moral beliefs; they are central to the function of moral cognition in human life. The moral grammarians, I shall argue, have no good account of this role of explicit moral judgement. Having established this prima-facie contrast between moral and linguistic cognition, I then deconstruct the moral version of the poverty of the stimulus argument, developing an alternative model of the development of moral cognition. That account combines three ingredients: pattern recognition; engineered learning environments, and prosocial emotions.

III

Reflective Moral Thought

Suppose the nativists are right. There is an abstract set of normative general principles that develop in every normal individual. Once switched on and initialised by experience, these principles automatically, rapidly, productively and

unconsciously generate moral appraisals. Even if all this is true, there is still a striking difference between language and moral cognition. For in contrast to language, every normal agent also has conscious, and somewhat reflective, moral principles. We endorse moral generalisations; we do not just make bullet-fast judgements about specific cases. Our culture provides us with a rich and nuanced set of tools for the moral evaluation of acts and agents; 'cruel", "capricious, "spiteful", "kind", "vengeful" and so on. Law has developed a technical vocabulary and so has moral philosophy. But those vocabularies are built on, and continuous with, the quite rich resources available to non-specialist agents. Thus once the doctrine of double effect is explained, many agents recognise it as the principle they judge by. No similar experience occurs when reading, say, the cognitive science of depth perception. For the technical vocabularies of linguistics, psychology or economics are not just developed versions of folk resources. What is the role of reflective morality on the Hauser-Mikhail-Dywer picture?

One possibility is that reflective morality is incomplete introspective access to the moral grammar. After all, experimental work seems to show that agents have some conscious access to the principles that underlie their judgements, For example, even though they cannot articulate the double effect principle, agents can articulate a principle appealing to the

moral difference between acts and omissions. Moral dumbfounding is not complete. However, reflective morality is unstable in individual agents, and variable across agents. Agents convert to utilitarianism; rejecting the act/omission distinction and the doctrine of double affect. They decide that mixed marriages or infertile incest are morally blameless. In contrast, once developed, a moral grammar is presumably fixed. Moreover, we do not always endorse or act upon our fast intuitive judgements. So if we have a moral grammar, conscious, self-aware moral thinking is a second system, not partial access to that moral module. We have both tacit and explicit moral cognition. How might the two systems interact?

Perhaps the moral grammarians do not need a model of interaction because the systems do not interact; perhaps reflective morality is mostly unconnected from agents' actual practices or moral evaluation. After all, Jon Haidt and his colleagues have argued that conscious moral reasoning is often the post-hoc rationalisation of rapid, emotionally mediated responses. But while no doubt conscious moral reasoning is sometimes confabulation, it would be very implausible to claim that reflective morality is epiphenomenal, and Haidt and co make no such claim. Agents change moral practices (converting to vegetarianism, for instance) as a result of newly acquired reflective moralities. Moral argument sometimes results in

agents changing their views, the ways they act, and (eventually) their fast-response intuitions. A key form of moral argument is to try to induce another to see similarity relations between cases they already evaluate morally in a distinctive way and the case under dispute: to bring intuitive fast judgement into interaction with reflective principles. And these arguments sometimes work. If there are two systems, they interact.

Indeed, the phenomenon of moral argument is itself an indicator of a pivotal difference between language and moral cognition. It is no accident that we have causally salient and explicit moral beliefs. Moral cognition does not just guide our own action; it plays a central role in persuasion, and that explains a crucial difference between linguistic and moral cognition. It is not just general principles of language which are tacit, unavailable to introspection. Structural representations of particular utterances are also introspectively unavailable. In hearing a sentence, in some sense we must compute its organization, for that organization plays a central role in sentence meaning. We understand sentences and to understand them we must represent sentence structure. But in the standard case, we do not have syntactic beliefs. The parsing system takes as input sound from speech, and gives as output a representation of utterance structure. But that output does not consist of beliefs about sentence structure. Few agents have any such beliefs. In

contrast, the output of a normative judgement system (if there is one) consists of beliefs about what should or should not be done. We may not always be able to articulate the general principles that guide our specific normative judgements. But we can certainly articulate those judgements themselves and make them public. We do not just find ourselves inclined to act one way rather than another.

Moral argument is a normal part of human social life because one role of moral cognition is to influence others' actions, not just guide one's own. To influence others, I must articulate my views and my reasons. Thus an account of explicit moral thought is central to any explanation of the evolution and role of moral cognition. We seek to persuade others to share our moral views both in particular cases and with respect to general principles. As Haidt and his colleagues note, moral reasoning is often collective and social. It is a central part of social life. Moralising is not a private vice. In making normative appraisals, agents are not just in the business of guiding their own behaviour; they are in the business of guiding and evaluating the actions of others.

The contrast with language is striking. There is no syntactic analogue of moral persuasion. So in contrast to syntactic judgement, moral judgement is in the public domain. We have

no interest in the syntactic judgements of others: the more invisible syntactic processing is to most users of language, the better. It is like our retinal image. Artists excepted, we do not want to know about retinal images. Rather, we want to, and we normally do, "see through" the image to the world. Similarly, we normally "listen through" the specific form of others' utterances, using our conversational partners as instruments that tell us about our common world. In conversation, we form beliefs about our environment not utterance structure. But we have very good reasons to identity the normative judgements of others. They are important indicators of how those others will act. And when we disagree, those views are the targets of our persuasive efforts. The role of moral judgement in human social life requires specific moral judgement (and perhaps the principles which guide those judgements) to be in the public domain. The function of moral cognition, then, implies that there will be explicit moral thought, and it is not epiphenomenal.

IV

Moral Pattern Recognition

So I do not think moral thought is modular. But moral judgement is fast, automatic and develops reliably. Those sceptical of the moral grammarians' picture will need to explain

these characteristics of moral thought. I begin that project with a discussion of pattern recognition. Natural minds are good at learning to recognise patterns — similarities between instances. The exercise of that capacity results in intuitive judgements about new cases. Pattern recognition is fast and automatic. Often the agent cannot explain the basis of their judgement. In a series of vivid examples, Chomsky argued that grammar could not be captured as set of simple lexical patterns. To capture (say) the relationship between indicative and yes-no questions, we need to identify the abstract structure — the organization of lexical items into subsentential constituents. If these arguments are sound, agents cannot not learn their language by generalising from familiar cases; using familiar examples as models for new constructions.

In my view, no-one has given a Chomsky-style argument to show moral learning is not largely pattern recognition. For while Marc Hauser and John Mikhail have argued that moral judgements depend on subtle situational facts, those facts are not moral but intentional. They are facts about goals, intentions, consequences. So one possibility is that our intuitive moral judgements are generalisations from exemplars. Kind and generous acts are those that resemble paradigmatic moments of kindness or generosity, and so on for other evaluations. Noticing and estimating these similarities may be cognitively complex,

and depend on subtle general principles. But these will be the principles of folk psychology, not deontological morality. Pattern recognition is fast and automatic, once the abilities have come on line. Chess experts, for example, assess positions rapidly and accurately, and it is very hard for a chess player not to see a chess position as a position. Moreover the metrics underlying pattern recognition are often tacit. Expert birders, for example, can often recognise a bird from a fleeting glimpse without being able to say how they recognise it. It has the "jizz", they will say, of a brown falcon. So if intuitive moral judgements are the result of pattern recognition capacities, it is no surprise that they have the rapidity and introspective opacity that Hauser and his colleagues have identified. If agents project to new cases systematically from their learning set, while lacking, as they often do, introspective access to their own similarity measures, we will find moral dumbfounding.

Furthermore, an exemplar-based view of moral judgement is independently plausible. For moral education is often largely-example based. Children are exposed to a rich stock of exemplars. The narrative life of a community — the stock of stories, songs, myths, and tales to which children are exposed — is full of information about what actions are to be admired, and which are to be deplored. Young children's stories include many moral fables. They are read stories of right action and

motivation rewarded; of vice punished. Their narrative world is richly populated with moral examples. So too, for many, is their world of individual experience. Children do not just look and listen. They act. In their interactions with peers, they encounter many morally charged situations, especially those to do with harms and with fairness. Children have many experiences annotated with their moral status to act as input to a pattern-recognition learning system.

Moreover, an exemplar-based view of moral intuition makes sense of an otherwise surprising fact: we judge harms by inaction less severely than harms of direct action. From an adaptationist perspective, the act/omission distinction is puzzling and arbitrary. My defection by inaction is just as costly to my partner or my group as a defecting act would be. Failing to help a partner in dire need dooms them as certainly as a malicious act would. If the adapted function of morality is to support and extend prosocial behaviour, we would expect omissions and commissions to be morally equivalent. Moreover, omissions are more difficult to detect. Passive deception — allowing you to act on a misapprehension that I recognise and could correct — will typically be more difficult to identify than a deliberate lie. The same is true of failures to provide material aid. It will often not be obvious to others whether the defecting agent was aware of the situation and in a position to intervene.

This intensifies the puzzle. Whether it is in my interest to defect depends on the reward of successful defection, the risk that the defection will be detected together with the severity of punishment, if detected. To keep a constant level of deterrence, as the risk of detection goes down, the severity of the punishment should increase. The well-adapted intuitive moraliser should be incandescent with rage at cryptic harming by omission, but somewhat more forgiving of upfront, in your face, fuck-you defection. The puzzle abates if intuitive judgement depends on generalisation from a set of exemplars. For the detection problem above will bias the training set in favour of acts of commission. The obvious, unmistakable examples of kindness or cruelty; of fairness or greed, will be things we do, not things we fail to do. Our similarity measures coalesce around a core of positive interventions. We judge omissions less harshly because they are further from our paradigms of the punishable.

In brief, I think it is likely that agents are unable to make explicit the basis of many of their moral judgements because those judgements are based on projection from exemplars rather than because their judgements are based on introspectively hidden, built-in abstract principles.

V.
Constructing The Moral Niche

I agree with the nativists in thinking that acquiring moral cognition is biologically prepared. We are adapted for moral thought in such a way that its development is accelerated and made more reliable. Moral development is robust because we are biologically prepared for moral education. But I think that preparation consists in the organization of our developmental environment and through specific perceptual sensitivity, rather than through pre-wiring tacit, general, abstract moral principles.

Consider, first, the organization of development. Parents do not just wait and hope that children will acquire the information they need. They organise the informational world of their children. They provide informational resources: toys, games and other props. The narrative experience of children is mostly provided by adults. Moreover, language itself is a powerful resource. It marks for children the similarities and differences the community takes to be important. And, of course, there is a significant amount of explicit teaching. Perhaps in contrast to language, children are supplied with explicit information about what is forbidden, not just what is permitted. In short, the

parental generation actively engineers the learning environment of the next generation.

Moreover, children themselves are far from passive: they are themselves active epistemic agents. Children do not acquire information about the moral opinions of their community just by observation of adult practice: of what adults do and avoid doing. Children experiment with and manipulate their surrounds. There is trial and error moral learning, but in a supervised environment: adults sometimes intervene in children's moral disputes, but often they do not. So children's social worlds are full of disputed terrain, especially to do with issues of fair division and responsibility. Children collide with the moral views of their peers, and they attempt to impose their own views on those peers. Few children could learn from the norms of fair division from simple induction on the actions of their brothers and sisters. But they have a good chance of learning them from overhearing and taking part in discussions over how spoils are to be divided; "you cut; I choose" and similar division rules. So a child's moral development normally takes place in a prepared environment.

VI

Moral Cognition and Prosocial Emotions

I have just argued that the downstream engineered of developmental environments is important. But we are also prepared for moral education by the phenomena we find perceptually salient. I agree with Jesse Prinz, Jon Haidt, Shaun Nichols and others who think that the moral grammarians place insufficient weight on the role of moral emotions in moral cognition. The task of the child is not just to discover the set of practices governing a particular community. Rather, her task is to join her community; to share rather than describe those norms. She can do so because moral cognition is embodied. We are perceptually tuned to the emotions and emotional responses of others, and to our own emotional responses to others. We typically notice other's distress, and we do not just notice it, we respond with emotions of our own. Those emotions, too, we notice by an internal analogue of perception. We respond positively to kindness; we are aware of our own positive response, and we convert that visceral reaction into a normative judgement. Moral norms are grafted on top of our dispositions to respond emotionally in characteristic ways to stereotypic stimuli. Normal children, for example, notice when their playmates are distressed. Their own emotional responses to the

emotions they notice are motivating. Distress makes most children uncomfortable. These emotional responses normatively mark certain types of events. As Nichols sees it, putative norms which line-up with these characteristic responses are more stable. They are more likely to be endorsed, respected and taught that arbitrary norms.

In short, our emotions make certain types of action and situation salient: we notice the emotions and reactions of others, and that in itself will narrow moral search space: in thinking about the content of moral norms, they will tend to think of human emotions, reactions, and the stimuli that engage them. Children (and to adults) notice these phenomena, and are motived by them, independently of, and prior to moral education. So a child's evidential base includes not just patterns in others' actions; it includes information about others' emotional responses to those actions, and information about that agent's own emotional response. So we are biologically prepared to develop moral cognition because moral cognition is a natural development of our existing emotional, intellectual and social repertoire. Our suite of emotional reactions — especially those concerned with reciprocation, sympathy, empathy disgust and esteem — shape and constrain moral cognition. Moral values are adapted to pre-existing modes of human cognition; we did not have to change to make norms comprehensible.

While I am sympathetic to those who emphasis the role of emotions, it is important to recognise a role for explicit principles. Those who emphasise the role of emotions sometimes underplay the role of moral reasoning. For example, Haidt and his co-workers take intuitive moral judgement to be the reflection of our emotional responses of aversion and approval. Conscious moral reasoning, in turn, is typically a post-hoc response to intuitive judgement. While not claiming that moral reasoning is always epiphenomenal, their picture is very bottom up. In contrast, in my view moral cognition develops from an interaction between emotions, exemplar-guided intuitions and explicit principles in richly structured environments. Haidt's work focuses on synchronic response. I think moral reasoning has its effects diachronically. So I think Marc Hauser is right to emphasise examples like moral vegetarians. As he points out, the disgust such agents come to feel for meat is a consequence rather than a cause of their moral convictions. While emotion is certainly intimately linked to moral appraisal, sometimes the appraisals cause the emotions. Moral vegetarianism and similar examples show that moral emotions and moral principles interact, and that over time principles can shape emotional responses. Top-down reasoning effects emotions, choice of exemplars, and, especially, generalisation from exemplars. Moral vegetarianism is, after all,

based on the idea that animals are relevantly similar to paradigm exemplars of morally valuable agents.

VI
Where Next?

The picture presented accepts that the development of moral cognition is supported by our evolved biology, but it does not invoke domain-specific principles on the model of grammar. Many human competences develop reliably in the downstream generation because the upstream generation structures downstream learning environments to support those developments. The development of a cognitive competence is often stabilised by ecological engineering, rather than a genetically-based special purpose cognitive mechanism. And it often depends as well on perceptual biases of various kinds. This general view may well be true of moral competence: it is transmitted culturally yet reliably. The parental generation engineers the informational environment in which the next generation develops, thus guaranteeing the development of moral competence.

This model is especially plausible if the informational environment is not stable, and I suspect that moral nativists somewhat overstate the cross-cultural similarities across

normative environments. Some normative systems are much more focused on respect for social hierarchy or on purity than on prosocial regulation and harm minimisation. Food taboos; obsessions with caste; phobias about menstrual fluids are all hard to see as a special cases of universal deontological maxims. The existence of norms of some kind is plausibly universal. The existence of specifically moral norms, let alone quasi-Kantian ones, may not be. The normative environment might be quite heterogeneous; none the less, children join their normative world, because the path is prepared for them by the previous generation.

I mentioned, though, in session 1 that it is possible to combine elements from the standard and alternative models to produce hybrids, and that is certainly possible here. Suppose, for example, I am right to suggest that children develop their moral intuitions about fairness on the basis of prototypical representations of fair and unfair actions. If a child has the concept of a moral transgression, perhaps it is not hard to work out which acts count as transgressions, given a suitable learning set. But perhaps acquiring that core concept is the crucial challenge to a non-nativist view of morality. Richard Joyce argues that the concept of a moral norm itself poses the key learning problem. He wonders how a generalized learning mechanism could develop the idea of a moral transgression,

even in a rich and varied environment. I am not convinced. But in the absence of an explicit account of moral concepts and of what is required to learn them, it would be premature to reject this minimal nativism. Importantly, though, such a minimal nativism is not just compatible with most of what I have said here; minimal nativism depends on rich, supported learning. It is one possible hybrid model.

Let me conclude this session, and introduce the themes of the next, by returning to adaptation and adaptationism. I rather causally and in passing accepted the idea that normative cognition is probably a co-operation-enhancing adaptation to human social worlds. That can be true even if, as suggested here, the development of moral cognition depend on informationally engineered learning environments. What, though, of specific systems of moral thought? Might they be adaptations to the specific social milieu in which we find them? Perhaps the differences I noted earlier reflect selection, albeit selection on culturally transmitted phenotype differences, rather than cultural contingency. For if norms flow with high fidelity and reliably across the generations; if they vary, and if they contribute differentially to high-stakes decision making, selection on moral phenotypes might adapt agents differently to their different environments. This general issue will be one focus of the next two sessions. So for now, I will just sketch the

reasons for suspecting that selection will not optimise our normative profile.

First, the power of selection in part depends on the supply of variation, and normative variation may be limited in two important ways. First, specific norms — of marriage; of childcare practices; kinship systems; property rights — do not vary independently of others. Specific norms tend to fuse into large norm complexes. To the extent that this is so, selection cannot optimise, say, norms of property rights. For they will not vary independently of, say, kinship systems. Second, as we saw earlier, our moral emotions make some similarities salient, easy to notice and remember. Those same emotions make other similarities much less salient or apparent. Hence norms based on those will, as Nichols points out, be much less likely to appear or establish. The more developmentally stable we take moral emotions to be, and the more important we take them to be in the early acquisition of norms, the more powerful they will be in constraining the supply of selectable variation in norms.

Second, technological practices are often matters of individual decision. I can adopt a new style of fish-hook making independently of others, and if I then transmit that to my own family, selection at the level of individual agents can respond to fitness differences so created. As norms are to a considerable

extent co-ordination devices whose function depends on their being shared, it is much less clear that individuals within a group or culture can create and transmit normative variation by individual decision. If not, the unit of variation and selection will be whole groups or cultures, and the mechanism of adaptation will be slower and less powerful. Selection is a much more powerful optimiser of fishhooks than it is of marriage practices.

Third, it may be that the fidelity of norm transmission is not high enough for optimising evolution. It is high enough for ethnographers to recognise a phylogenetic signal: we can trace honour cultures through their remarkably persistent norms. But optimisation requires cumulative adaptation, and that in turn makes small innovations important. These must be retained at a generation once made, and reliably transmitted to the next. The very mechanisms of conformity that tend to make groups normatively uniform, and which add redundancy and reliability to cross-generation information flow, will tend to work against the cumulative improvement of normative systems.

Questions about the evolution of information sharing, and the differing evolutionary and cognitive aspects of different mechanisms of information sharing, are the focus of the next session. The take-home message of today is that moral nativism

is not mandatory. Normative cognition could be genuinely universal, an adaptation that develops robustly without that development relying on innate, tacit, abstract principles.