

Automaticity in Virtuous Action

Clea F. Rees and Jonathan Webber

The Philosophy and Psychology of Virtue, ed. Nancy Snow and Franco Trivigno
Routledge, 2014

Penultimate draft. Please cite only published version.

Automaticity is rapid and effortless cognition that operates without conscious awareness or deliberative control. It is the subject of much empirical research in contemporary psychology. It has traditionally been emphasised in discussions of ethical virtue. Our focus in this paper is on the relation between automaticity and virtuous action. An action is virtuous to the degree that it meets the requirements of the ethical virtues in the circumstances. What contribution does automaticity make to the ethical virtue of an action? How far is the automaticity discussed by virtue ethicists consonant with, or even supported by, the findings of empirical psychology?

We argue that it is a mistake to apply the analogy between virtue and skill here. The automaticity of virtuous action is automaticity not of technique but of motivation. Skillful action can be admirable and can indicate the presence of ethical virtue, but does not itself contribute to the ethical virtue of an action. The automatic motivations that do contribute to that virtue are not only those that initiate action. They can also be motivations that modify action that is otherwise initiated, or that initiate and shape practical deliberation.

We then argue that both goal psychology and attitude psychology can provide the cognitive architecture of this automatic motivation, although goal psychology is not as advanced as attitude psychology in modelling the habituation traditionally thought to be involved in the acquisition of ethical virtue. Moreover, since goals are essentially directed towards the agent's own action whereas attitudes are not, we argue that goals might underpin some virtues while attitudes underpin others. We conclude that consideration of the cognitive architecture of ethical virtue ought to engage with both areas of empirical psychology and should be careful to distinguish among ethical virtues.

1. Automaticity in Two Rescues

Philippa Foot argues that an individual's virtue is a matter of 'innermost desires' as well as intentions, of 'what is wished for as well as what is sought' (1978: 5). The example Foot uses to illustrate this point suggests that 'innermost desires' are essential not only to the virtue of an agent, but also to the virtue of action. This is the example of a tracker rescuing an injured boy from a river. The tracker hears the distressed boy and immediately responds in the right way to rescue effectively. The automaticity of this rescue is part of the virtue of the action, in that it would seem not merely less likely to be effective but also less virtuous had the rescue involved deliberation.

The tracker is capable of responding without deliberation as a result of training himself to notice when someone is in danger and to respond accordingly. The automaticity, that is to say, has come about through habituation. It might be argued that for action to be truly virtuous, the dispositions it manifests must have come about through such habituation. Dispositions that have come about by heredity or accident, on such a view, could not be virtuous however good the actions that manifest them might be. We will not consider this issue in this paper.

Our concern is restricted to automaticity itself, rather than its origins. There are several kinds of automaticity in this case. Which of these are essential to the ethical virtue of the action? To begin to separate out the different kinds of automaticity involved, we can compare this rescue with another described by Bernard Williams. In this case, the rescuer is able to rescue only one of two imperilled people, one of whom is the rescuer's wife. Williams argues that if the rescuer thinks 'that's my wife, and in situations of this kind it is permissible to save one's wife', then the rescuer has had 'one thought too many', since it would be better if 'the motivating thought, fully spelled out' were 'the thought that it was his wife' (1976: 214-5). The thought about permissibility does indeed seem out of place. But being motivated by the thought 'that's my wife' also seems less than virtuous. A more virtuous action would not be motivated by the fact of being married to this imperilled person, but would be motivated by a concern for that individual's welfare sufficiently central to the agent's outlook to be automatic.

This rescuer may need to think about how to rescue the drowning wife. Unlike the tracker, this rescuer may have no experience of dealing with dangerous waterways. But this is distinct from having to think of a reason to rescue her. Moreover, the two rescuers differ not only in their experience of dangerous situations, but also in the content of the motivation that is automatically activated. The tracker's motivation is to ensure safety in dangerous circumstances. The other rescuer's is to care for this particular individual. Each of these motivations leads to action whose immediate goal is to rescue an individual, even though neither motivation specifies that action in its content.

A second difference between the two rescuers concerns their skill in rescuing. The tracker displays great skill where the other rescuer might not. Again, this is due to the difference in their background experience. Does this mean that the tracker's rescue is the more virtuous? It is overall a more excellent action, but it does not follow that it is more ethically virtuous. For it does not seem that the ethical virtues demand that everyone be a skilled

rescuer. This might be a requirement for people whose lives involve regularly dealing with dangerous waterways, so perhaps such skill is required of the tracker. But assuming our other rescuer lives a very different kind of life, the lack of skill involved in this rescue does not detract from its virtue.

This is not to say that fully virtuous action requires only its initiating motivation to be automatic. We will argue that virtuous action can be a matter of the automatic guidance of action, but that this is a matter of modifying motivations rather than skill. In the next three sections, we consider in more detail the automaticity of skilled action, of modifying motivations, and of initiating motivations, before going on to consider the cognitive architecture that can underpin automatic motivation in virtuous action.

2. Skill and Action Automaticity

Recent discussions of the structure of virtue have placed significant emphasis on the idea of skill. There are important analogies between virtues and skills. In particular, both are the product of habituation through rationally guided practice rather than acquirable through propositional learning alone (Annas 2011: 16-32). More strongly, it can be argued of either that its content cannot be wholly captured in propositional form.

However, there are important differences too. A skill is an ability or capacity, whereas a virtue essentially involves a tendency. Many more people have the ability to behave honestly than possess the virtue of honesty. Likewise, knowing how to respond courageously or understanding what fairness requires are not virtuous in the absence of the motivation to behave accordingly. One cannot become virtuous just by learning any particular ability, since one must also come to have the right motivations. In the absence of the right motivations, any skill can form part of a vice rather than a virtue (Ryle 1972: 438-41).

Moreover, the contrast between the two rescue cases shows that, at least in some instances, virtuous behaviour does not require any skill at all. Our two rescues are equally ethically virtuous, though only one of them exhibits skill in rescuing. These cases also indicate that an action is more virtuous if its motivation is automatic rather than deliberative. In the terminology employed by Bill Pollard and Nancy Snow, our rescues are fully ethically virtuous if there is no deliberation about whether to rescue, regardless of whether there is any skill involved and regardless of whether there is deliberation about how to rescue.¹

Is this a general point about virtue? Or do some virtues essentially involve skills? The skill present in one rescue case but not the other is not itself a matter of ethical character, of values and commitments, but only of physical ability. Peter Goldie's example of a dinner

¹ Our rescue cases thus conform to Pollard's principle that to be fully virtuous an action cannot involve deliberation about whether to act, but may involve deliberation about how to act (2003: 416-7). Snow raises important objections to this as a general principle, but retains the distinction between whether and how to act (2010: 48-9). In this paper, we argue that this distinction between whether and how to act should be replaced with a distinction between initiating and modifying motivations.

party conversation might seem different in this regard. As one of the diners is becoming upset by the conversation, another diner notices this and discreetly changes the subject. Goldie argues that the virtue of this action does not consist solely in the motivation to save the other diner's feelings, but also in noticing that she is becoming upset and in being sufficiently sensitive to know that changing the subject is the best way to save her feelings (2004: 44). Should we understand this noticing and this knowing how best to resolve the problem on the model of skills? Or are they aspects of automatised motivation?

There are delicate issues here. Noticing that someone is becoming upset manifests a concern for the feelings of that person. The noticing cannot be understood as automatically initiating that concern, since the noticing occurs only because the concern is already guiding the agent's cognition. This guidance is a way in which a motivation can automatically modify action. In this case, the motive of concern for the other person's feelings modifies the action of engaging in a conversation. The motive of concern must have its influence automatically rather than deliberately if it is to make a feature of the situation salient to behavioural cognition, since deliberation deals only with such features once they are salient.

Is knowing how best to save the other person's feelings a skill, or should it too be understood as automatically manifesting a modifying motivation? In this case, mastery of a technique seems essential. The person who clumsily comforts the other person, thereby drawing attention to their distress, and the person who adroitly changes the subject need not differ in the concern they express or in the automaticity of that concern. The difference seems rather to lie in knowing how best to achieve the aim of saving the other person's feelings. So this knowhow does seem to fit the model of a skill.

However, it is not clear that this skill is essential to the action being virtuous. The person who changes the subject seems more virtuous than the one who clumsily attempts to comfort. But it does not follow directly from this that this action of changing the subject is more virtuous than the action of comforting. For it might rather be that knowing how to achieve an aim is often indicative of having had experience of trying to achieve that aim. Knowing how to act on a motivation is thus indicative of having habituated that motivation. At least, it is indicative of having habituated that motivation in contexts relevantly similar to the present one. When the motivation has only been habituated in importantly different contexts, one might automatically be motivated in the right way but not know what to do. This is the difference between our two rescue cases.

In the dinner party case, the motivation of kindness towards others seems unlikely to have been habituated without having been acted on in situations where other people's feelings might be hurt. This is because situations in which other people's feelings might be hurt are pervasive in everyday life for most people, unlike situations in which someone needs to be rescued from a river. So it seems unlikely that anyone could have automatised the motivation of kindness without also developing the relevant social skills through the very same actions. This is not to say that it is impossible to habituate this motivation without developing the relevant skills. It is only to explain why social skills would usually accompany the virtue of kindness even if skill is not itself part of virtue. The person who changes the subject to spare their fellow diner's feelings probably is more virtuous than

the person who tries to comfort the fellow diner, but this is due to the difference in the degree to which they have habituated their concern for the feelings of others; it is not a difference in how well the act itself meets the demands of ethical virtue.

This suggests that the right automatic motivation might be all that is required for an action to be virtuous. The analogy between virtue and skill, therefore, should not be pressed too far. But we do need to distinguish between motivations that initiate actions and those that modify actions that are otherwise initiated. This is a refinement of the distinction between whether and how to act. It is required to explain how virtuous motivation can be responsible for the agent noticing relevant features of the situation. In the next two sections, we consider modifying and initiating motivations in more detail.

3. Automaticity of Modifying Motivations

John Bargh and colleagues have performed various experiments to demonstrate that motivations can automatically modify actions. One such experiment is their 'fishing game' (Bargh et al 2001: 1017-9). Participants play this game on a computer, believing themselves to be playing with another participant but in fact this is simulated by the computer. In each season, the participant catches fifteen fish and must decide how many to return to the lake. The number of fish returned determines how many new fish appear in the lake before the next season, according to the formula $5n-30$. Returning ten fish, for example, would mean that a further twenty appear in the lake. The lake begins with one hundred fish and if the number of fish in the lake falls below seventy then both participants will lose all the fish they have caught. One can play this game cooperatively by aiming to maximise fish for all players or play competitively by aiming to collect the most fish.

Participants entered this game in one of four conditions, on a two-by-two experimental design. On one dimension, participants were either explicitly instructed to play cooperatively or given no explicit instruction on the spirit in which to play the game. The other dimension concerned a scrambled sentence task that all participants completed before the game and were told was unconnected to the game. Participants had to make four-word sentences from lists of five words. For some participants, one-third of these lists of words contained terms relating to cooperation. For the other participants, there was no such pattern in the word lists.

Those whose scrambled sentence task had included words related to cooperation played significantly more cooperatively than did those whose scrambled sentence task did not include words related to cooperation. This was true for those given explicit instruction to cooperate and for those not given this explicit instruction. Similarly, for each scrambled sentence task, those given explicit instruction to cooperate did so to a significantly greater degree than those not given this explicit instruction.

Subtle cues in the scrambled sentence task therefore caused the action of playing the game to be modified by the motivation to cooperate. The explicit instruction to cooperate also had this effect. The most cooperation was found among those participants who had been given both the explicit instruction and the subtle cues. The influence of the explicit

instruction is not a case of automaticity. Participants consciously make an explicit decision in response to that instruction. Mostly, they decide to obey it. But the subtle cues operate below the radar of conscious awareness. Participants were unaware of this influence on their behaviour.

This influence is best understood as a case of priming, which is the causing of a temporary increase in the accessibility of a mental state. The more accessible a mental state is, the more readily and more rapidly it is brought to bear on cognition, so the greater influence it has over that cognition, whether this influence is conscious or unconscious. Each mental state has a chronic degree of accessibility, which is temporarily increased by the state's being activated either in the course of cognition or as a direct result of some relevant situational cue. In this instance, the motivation to cooperate has been unconsciously activated by the subtle cues in the scrambled sentence task, which has the effect of this motivation having a greater effect on the agent's cognition during the fishing game than it would have had otherwise.²

Chronic accessibility is itself the product of the mental state having been brought to bear on cognition in the past. The more a mental state has been employed in cognition, the higher its baseline level of chronic accessibility, so the greater its influence on cognition generally. Priming is merely the temporary raising of accessibility above this baseline. Somebody who had repeatedly worked at behaving cooperatively would have habituated their motivation of cooperativeness to a high degree of chronic accessibility and would thereby be likely to behave cooperatively in the fishing game even in the absence of explicit instruction or subtle cues encouraging them to do so.³

Because chronic accessibility is a matter of degree, so too is the modification that a habituated mental state can make to cognition. In the dinner party example, the virtuous person who notices their fellow diner's growing distress and without a thought changes the subject is someone who has habituated the motivation of kindness to a very high degree of chronic accessibility. This is why the agent not only notices the distress, but is sufficiently strongly motivated to alleviate it. The skill involved in alleviating it is not itself a matter of this habituated motivation, which could be equally chronically accessible in an agent who noticed the distress and was strongly motivated to alleviate it but did not know how to do so.

An agent who has not yet habituated their motivation of kindness to quite the same degree might notice the impending distress but need to deliberate about whether to intervene to

² Some psychologists, including Bargh, use the term 'priming' more narrowly, to mean a temporary increase in accessibility that occurs without the agent consciously noticing. Because there is no equivalent term for temporarily increasing accessibility in a way the agent does consciously notice, we use the term 'priming' more broadly to encompass any temporary increase in a mental state's accessibility.

³ Some writers use the term 'chronically accessible' as though it were binary rather than scalar. This is an unfortunate byproduct of the tendency to contrast a state that has a very high degree of chronic accessibility with one that has a very low degree of chronic accessibility by referring to the former as the 'chronically accessible' one. All mental states are chronically accessible. They differ in their degree of chronic accessibility.

alleviate it or to let the conversation take its course. In this case, the motivation is sufficiently highly accessible to bring the agent to notice the distress, but not accessible enough to automatically motivate action to alleviate it. An agent who had not habituated the motivation of kindness even to this degree could notice the growing distress only as a result of a deliberative decision to consciously look for signs of distress. Such a decision would be an early step on the road to a more habituated, more highly accessible, more automatic motivation of kindness.

4. Automaticity of Initiating Motivations

Automatic motivation does not only modify actions entered into as a matter of conscious decision, such as dinner party conversations and experimental fishing games. Actions can also be initiated automatically. Our two rescue cases are clear examples of automatic initiation of complex action. These are virtuous actions partly because the action of rescuing the imperilled person is initiated without deliberation, though there may be deliberation about how to effect the rescue. But there are also other roles that automatic initiation can play in virtuous action.

One such role is exemplified by the example of driving a familiar route home without consciously paying attention. This clearly involves skill automaticity. One major part of learning to drive is instilling motor routines concerning turning left, turning right, changing gear, and so forth. While it is indeed unlikely that the initiation of the overall complex action of driving home would ever itself be automatic rather than deliberative, we should not conclude that the skill involved in driving is the only automaticity in this case. For where the route driven is sufficiently familiar, many smaller actions that make up the overall action of driving home will be automatically initiated. Such actions as slowing down when approaching a set of traffic lights or indicating a left turn two minutes after joining a particular road at a particular junction could be automatic not only in the skill employed in their execution but also in their initiation in response to visual cues.

Such responses to visual cues should be understood as initiating motivations of actions that comprise driving home, rather than as modifying motivations that influence the overall action of driving home, precisely because they are responses to particular situations encountered along the way. They are not analogous to the motivations of kindness or cooperation that can influence the way you engage in a dinner party conversation or play the fishing game. Modifying motivations are active throughout the action they modify. Driving home might be modified in this way by such motivations as being safe and obeying the law. Moreover, such modifying motivations might be causally responsible for the automatic activation on seeing a red light of the motivation to slow down. But it remains that the motivation to slow down is an initiating motivation, one that initiates the action of slowing down rather than continually modifying the action of driving home.

A further kind of automatic initiation is well illustrated by Snow's example of the compassionate person whose friend has been made unemployed. Such a person, Snow argues, might fully possess the virtue of compassion, yet still need to deliberate over whether to respond to this news by offering financial assistance or whether instead to

remain quiet to preserve the friend's pride. Snow takes this case to show that virtuous action need not involve automaticity at all, since the action this person decides upon might be fully virtuous (2010: 48-9).

However, Snow has overlooked the question of the motivation for deliberating over whether to respond to the news. The deliberation itself is initiated by compassion for the friend. Is this initiation automatic or deliberative? It could be either, but it would be more virtuous if it were automatic. For this would be an immediate recognition that the situation requires a response that is best for the friend who has suffered the misfortune. This immediate recognition makes the deliberation about whether or not to offer assistance more virtuous than it would be if the agent had first decided through deliberation that they should do what is best for their friend.

This is an instance of action being virtuous despite involving deliberation over whether to act at all. But the resulting action might still have been virtuous had the agent just known immediately that this particular friend would not be offended by an offer of help, so had not needed to deliberate to reach the decision to offer help. By contrast, there might be actions that would not be virtuous unless brought about through deliberation. For example, we have argued in another paper that some moral dilemmas are such that no action in response to them can be virtuous unless it involves serious consideration of the problems with each option. Faced with the scenario in which Williams places his character Jim, for example, who must either kill one innocent person or allow twenty innocent people to be killed, the person of integrity will not find the answer easy or obvious, whichever way they decide to go (Rees and Webber forthcoming).

As with the case of deliberating over whether to offer to help the friend, this deliberation over whether to kill one innocent person or allow twenty innocent people to be killed will count as virtuous only if it is initiated by the right motivation. For if it is initiated by the motivation to ensure public approval, for example, or the motivation to act like a comic book hero, then it would not be virtuous. If the motive is to do what is right, then so long as certain other conditions are met it would be virtuous. Moreover, if this motivation is itself the outcome of deliberation about which ends one should try to serve in this situation, then the action would be less virtuous than if it is automatically initiated by the motivation to do the right thing.

An initiating motivation need not initiate action, therefore. We should recognise that virtuous action can arise from deliberation that is itself initiated automatically by a motivation that sets its goal, and indeed that some actions would be less virtuous if they did not come about in this way. We should also recognise the possible role of automaticity in initiating actions that are not virtuous in themselves, but which constitute a larger action that is virtuous.

5. Goals and Attitudes as Automatic Motivations

Automaticity in virtuous action concerns initiating and modifying motivations. This distinction has three advantages over the distinction between whether and how to act. First, modifying motivations account for the influence of a habituated virtue over the

agent's perception of situations. Kindness can modify one's dinner party behaviour, making one more likely to notice a fellow diner's growing distress. Second, there are cases where it is arbitrary to classify the relevant automaticity as concerning whether or how. In the fishing game, there seems no reason to classify the priming as influencing how the participant played rather than whether they cooperated. Likewise, there seems no reason to classify an automatic action that forms part of a larger action, such as slowing down when approaching traffic lights as part of the action of driving home, as cases of whether to do the smaller action rather than how to do the larger one. Third, the automatic initiation of deliberation is neither a case of whether to act nor a case of how to act. Responding to a friend's bad news by deliberating about whether to offer help is virtuous when that deliberation has the goal of doing what is best for the friend and has this goal set automatically.

In the terms of our distinction, a modifying motivation is one that is continually active throughout the action it modifies. This is why it can influence the agent's perception of the situation their action engages with. The motive of cooperation modifies the playing of the fishing game because it is a continual influence. An initiating motivation need not begin an action, since it might instead initiate deliberation concerning some particular end. When it does begin an action, the action might be part of a larger action. The motive to slow down when approaching traffic lights initiates the slowing down at the relevant point. It is not a modifier of the overall action of driving, even though it might itself be prompted by the modifying motivation of driving safely. The same motivation, moreover, might modify one action and initiate another as a result. The motive of kindness might modify an agent's dinner party conversation, so that the diner notices the growing distress of another diner, with the result that the same motive of kindness initiates the action of changing the subject. Likewise, a motivation might modify an action that it also initiated. The tracker's concern for the safety of others might modify the rescue that it initiated.

This distinction between initiation and modification, then, is a distinction between psychological roles that a particular motivation can play. It does not rule out a particular motivation playing one role in some cases, the other in other cases. One individual's kindness, for example, can initiate some of their deliberations, initiate some of their actions, and modify some of their actions. Indeed, as we have seen, it can modify one action with the result that it initiates another action which it then modifies. Since these are clearly distinct psychological roles, a motivation can play either role independently of whether it plays the other with respect to the same action. There are no cases, therefore, where it is simply arbitrary whether to describe it in one way or the other.

If this account of virtuous motivation is to be psychologically respectable, if there is to be reason to believe that such forms of automaticity are achievable by human agents, then the account needs to be grounded in an empirically supported theory of the nature and development of automatic motivation. Experimental psychology does not provide any such complete theory. But it does contain two areas of research that seem promising.

One of these is research into goal automaticity, which Snow recommends as a cognitive architecture of virtuous automaticity (2010: ch. 2). Bargh is the leading researcher into goal automaticity. He understands a goal to be a mental state that motivates a specific kind

of action and structures experience accordingly. Somebody viewing a house with the goal of burgling it will notice and remember different features than would be noticed and remembered by someone viewing the same house with the goal of buying a home (Bargh 1990: 97).

Bargh was originally concerned with a goal's automatic influence over cognition in response to features of the environment. He argued that such automaticity can result from a goal's association with an overall situation, such as a dinner party, or the goal's association with a more abstract social interaction, such as somebody becoming distressed (Bargh 1990: 114). Bargh developed this theory by providing evidence that goals persist and continue to influence cognition even in the absence of any conscious awareness of them or any environmental cues with which they are associated. Once a goal has been made sufficiently accessible, its automatic influence is no longer reliant on situational priming (Bargh et al 2001: 1019-21, 1023-4).

A second candidate empirical grounding for this account of virtuous motivation is provided by attitude psychology. This is a research tradition that is nearly a century old. Its most famous strand is cognitive dissonance theory, which focuses on the changes in an agent's attitudes in response to their own behaviour. But the tradition is by no means limited to this and has included work by many psychologists of varying methodologies and interests. Although this makes it difficult to identify a single detailed theory resulting from the research, there is certainly a core conception of attitude that emerges from it. This is the conception of a cluster of cognitive and affective mental states that together constitute the agent's overall evaluation of some object. (For overviews of attitude psychology, see: Cooper 2007; Fazio and Olson 2007; Maio and Haddock 2010.)

An attitude's object can be at any level of abstraction, so that you can have an attitude towards something as abstract as democracy or as concrete as the British parliamentary electoral system, as abstract as free jazz or as concrete as your particular copy of Albert Ayler's *Spiritual Unity*. The attitude's content is constituted by the contents of the mental states that make it up. An overall positive attitude to democracy, for example, might be made up of such items as a belief that democracy is the best way to keep the peace, a belief that it is the only political system that respects the autonomy of the citizens, a desire that peace be kept, a desire that autonomy be respected, and so on. An attitude is not necessarily a persisting mental state in its own right. It can be constructed from relevant mental states when an overall evaluation of the object is required. It need not subsequently persist in any form other than that set of distinct mental states.

For an attitude to persist as a mental state in its own right requires that the constituent mental items are bound together by strong associative connections. This ensures that the attitude as a whole is brought to bear on cognition whenever any of its constituent parts are brought to bear. Such an attitude functions in cognition as a whole, rather than as a set of disparate items. In the terminology of attitude psychology, the strengths of the associative connections between the constituents determine the overall strength of the attitude. This degree of influence an attitude has over cognition is determined by the rapidity with which the attitude is brought to bear on cognition, its degree of chronic accessibility, which itself is determined by its strength.

One of us has argued that this conception of attitude provides an appropriate cognitive structure for virtue ethics (Webber 2013). As we have seen, the same has been claimed for the conception of goal that can be drawn from Bargh's research. Should virtue ethicists prefer either one of these or should they embrace both?

6. Goals and Attitudes in Ethical Virtue

Automaticity is not the only important feature of ethical virtue that needs to be reflected in its cognitive architecture. Ethical virtues are dispositions that are characterised in terms of the behavioural outcome that they incline the agent towards, but this inclination is not usually understood to be restricted to a specifiable set of circumstances. And virtues are traditionally understood to be dispositions that can be gained through habituation, which means that they are both strengthened and refined through reflective practice. We will consider how goals and attitudes can underpin these aspects of ethical virtue before going on to consider the differences between goals and attitudes.

That ethical virtues are generally tendencies towards particular kinds of outcome, rather than traits that dispose the agent to respond to a particular kind of situation in a particular kind of way, has been somewhat overlooked in the recent debate over whether the idea of ethical virtue is consistent with the findings of contemporary psychology. The focus on experiments that measure the behavioural impact of altering features of the situation has brought with it a focus on the reactive aspect of ethical virtue. Ethical virtue is not simply a matter of one's reactions to particular circumstances, however, but also guides the kinds of situations one tries to bring about irrespective of the situation one is starting from, determines the kinds and degree of situational detail one takes into account in deliberation, and governs one's attitudes towards one's own dispositions and those of others (see Webber 2013: § 2).

Both goals and attitudes can provide the cognitive architecture of personal dispositions of this kind. The idea of chronic accessibility is central to both goal psychology and attitude psychology. The more accessible a psychological item is, the greater its influence on cognition generally. The item does not need to be directly related to a feature of the situation to exert influence over the cognitive processes going on in that situation. What matters is how accessible that item is to each cognitive process. An item with a high degree of chronic accessibility is one that is continually highly accessible, and thus one that influences cognition generally rather than only in response to particular features of situations. This general influence can therefore account for the influence a character trait has over the agent's cognition that is not simply determined by the features of situations they are in.

Moreover, both goals and attitudes increase in chronic accessibility each time they are brought to bear on cognition. A goal or attitude is strengthened whenever it is automatically activated and whenever it is called to mind in deliberation. The stronger it gets, the greater its accessibility. In this respect, both goals and attitudes fit one part of the idea that one becomes virtuous through reflective practice. By acting on a motivation and

by employing it in deliberation, one increases the influence of that motivation over subsequent cognition and hence behaviour.

Habituation involves more than just strengthening the influence of a particular motivation, however. Aristotle emphasises the role of reflective practice in refining the content of a motivation. The sensitivity to situational detail characteristic of ethical virtue requires that it is developed through engaging with those details, rather than simply learned as a rule to follow. Attitude psychology fits this aspect of habituation very well. Since an attitude is a cluster of other mental states, it is continually refined as it is considered and applied in new situations. New beliefs and desires are brought into the cluster that determine the precise content of the attitude (Webber 2013: § 4).

Goal psychology is less clear on this aspect of habituation. There is much suggestive material in Bargh's papers concerning the relations between goals and situations and concerning the relations between goals and behavioural plans. But this has not yet been developed, by Bargh himself or by other goal psychologists, into an account of the cognitive architecture of goals. In the absence of such an architecture, we cannot say with confidence whether goal psychology could underpin the role of habituation in refining the content of a virtue.

It might seem that this limitation of goal research provides good grounds for virtue ethicists to draw exclusively on attitude psychology for the cognitive architecture of ethical virtue. However, this would overlook an important difference between attitudes and goals. Attitudes are overall evaluations of their objects. Any implications they have for behaviour are the result of their influence over behavioural cognition. They do not themselves specify that the agent should try to achieve anything in particular. Goals, on the other hand, are centrally concerned with the agent's own actions. An individual's positive attitude towards democracy, for example, does not have as its object that individual doing anything in particular, even if it does explain why they regularly vote in elections. It can be satisfied by other people bringing about or upholding democracy. An agent's goal of being a good parent, on the other hand, is itself directed at behaviour of their own, however abstractly specified. It cannot be satisfied by someone else being a good parent.

This reflects a difference among ethical virtues. To possess the virtue of compassion, for example, one must be strongly averse to other people suffering and in favour of such suffering being alleviated when it occurs. If one's concern for this is sufficiently strong, relative to one's other concerns, then one will regularly act on it. But if one is not concerned simply with the prevention or alleviation of other people's suffering but is rather concerned that it be oneself who prevents or alleviates that suffering, such that one would not be satisfied by someone else alleviating suffering that one could have alleviated oneself, then one's outlook seems not to be virtuous. Indeed, this would be a kind of narcissism. The proper object of compassion is the well being of other people, not the role of oneself in that well being. On the other hand, one's own action does seem central to some other virtues. Integrity, for example, is a strong concern with getting one's own actions right. A positive attitude towards people generally behaving well is not a form of integrity.

Thus it seems that neither attitudes nor goals alone could provide the cognitive architecture of ethical virtue. It seems rather that attitudes could provide the architecture of those virtues that are centrally concerned with states of the world irrespective of who brought those states about and goals could provide the architecture of those virtues centrally concerned with one's own action. If this is right, virtue ethicists should look to both attitude psychology and goal psychology for the cognitive architecture of ethical virtue.

7. Cognitive Architecture and Ethical Virtues

Our examples of automaticity have all been ones where the motivation might have been an attitude or might have been a goal. Rescues can be motivated by attitudes of concern for others or for a particular person, or by such a chronic goal as being a hero or being a good partner. Cooperative play in the fishing game could be motivated by a general attitude in favour of cooperation for mutual benefit or could be motivated by the goal of being cooperative. Noticing a fellow diner's growing distress and changing the subject could be motivated by an attitude of kindness or a goal of behaving kindly. Slowing down when approaching traffic lights could be motivated by a positive attitude towards safety or the goal of driving safely. Deliberating about whether to offer to help a friend could be motivated by a concern for that friend's well being or a goal of being a good friend.

The virtue of each of these actions might depend not only on whether the motivation was automatically initiated, but also on which kind of motivation it is. For some cases might be analogous to compassion, where a positive attitude towards the well being of others is essential to the virtue whereas similar behaviour motivated by the goal of being the person who offers help seems narcissistic. Other cases might be analogous to integrity, where it is essential to the virtue that one is concerned about one's own behaviour rather than about states of the world itself. For this reason, virtue ethicists interested in the cognitive architecture of virtue should be concerned with the psychology of both attitudes and goals.

It might be argued that attitudes and goals are not really very different. Why not think of goals simply as those attitudes whose objects are, or include, one's own behaviour? But it is not clear what would be gained by this. We would gain a collective term that ranges over the two kinds of motivation, but at the cost of losing the term that uniquely picks out one of these two kinds. Moreover, an important question would be obscured by this terminological shift. Attitude psychology has converged on a particular understanding of the structure of attitudes, as clusters of mental states held together by associative connections of varying strengths. Goal psychology has not yet formulated a clear structure of goals. Perhaps these will turn out to be structured in basically the same way as attitudes. Or perhaps they will not. Only further research will decide this, research that will include philosophical analysis as well as experimental work.

An alternative outcome might be that goals include attitudes as constituents. We have argued here that a cognitive architecture based on goal psychology might well fit the virtue of integrity, since this virtue seems essentially concerned with one's own behaviour.

In another paper, we have argued that the virtue of integrity can be understood as a specific collection of attitudes (Rees and Webber 2013). These two claims are compatible. It may be that the goal at the heart of integrity is itself constituted by that set of attitudes. If this were to be true of that goal, then it would not follow that all goals consist in sets of attitudes. But it would provide a model of the structure of at least some goals that might have been obscured by thinking of goals as a variety of attitudes.

The right understanding of the role of automaticity in virtuous action will rest in part on the outcomes of such empirical and philosophical investigation into cognitive psychology. But it will also rest on more careful consideration of the differences between the virtues, which the discussion of their relation to experimental psychology has tended to assume are uniform in this regard. The focus of this work ought to be on the automaticity of initiating and modifying motivation. It is not clear that skill has any defining role in the ethical virtue of action.

References

- Annas, Julia. 2011. *Intelligent Virtue*. Oxford: Oxford University Press.
- Bargh, John A. 1990. Auto-Motives: Preconscious Determinants of Social Interaction. In *Handbook of Motivation and Cognition: Foundations of Social Behavior* vol. 2, edited by Richard M. Sorrentino and E. Tory Higgins, 93-130. New York: Guilford Press.
- Bargh, John A., Peter M. Gollwitzer, Annette Lee-Chai, Kimberley Barndollar, and Roman Trötschel. 2001. The Automated Will: Nonconscious Activation and Pursuit of Behavioral Goals. *Journal of Personality and Social Psychology* 81 (6): 1014-1027.
- Cooper, Joel. 2007. *Cognitive Dissonance: Fifty Years of a Classic Theory*. London: Sage.
- Fazio, Russell H. and Michael A. Olson. 2007. Attitudes: Foundations, Functions, and Consequences. In *The Sage Handbook of Social Psychology: Student Edition*, edited by Michael A. Hogg and Joel Cooper, 123-145. London: Sage.
- Foot, Philippa. 1978. Virtues and Vices. In *Virtues and Vices and Other Essays in Moral Philosophy*, by Philippa Foot, 1-18. Berkeley: University of California Press.
- Maior, Gregory and Geoffrey Haddock. 2010. *The Psychology of Attitudes and Attitude Change*. London: Sage.
- Pollard, Bill. 2003. Can Virtuous Actions be Both Habitual and Rational? *Ethical Theory and Moral Practice* 6:411–425.
- Rees, Clea and Jonathan Webber. 2013. Constancy, Fidelity, and Integrity. In *The Handbook of Virtue Ethics*, edited by Stan van Hooft, pp-pp. Durham: Acumen.
- Rees, Clea and Jonathan Webber. forthcoming. [paper under review].
- Ryle, Gilbert. 1972. Can Virtue Be Taught? In *Education and the Development of Reason*, edited by Robert F. Dearden, Paul Heywood Hirst, and Richard Stanley Peters, 434-447. London: Routledge and Kegan Paul.
- Snow, Nancy. 2010. *Virtue as Social Intelligence: An Empirically Grounded Theory*. New York: Routledge.
- Webber, Jonathan. 2013. Character, Attitude and Disposition. *European Journal of Philosophy*. DOI: 10.1111/ejop.12028
- Williams, Bernard. 1976. Persons, Character and Morality. In *The Identities of Persons*, edited by A.O. Rorty, 197-216. Berkeley: University of California Press.