

CHAPTER 5

*Practical competence and fluent agency**

Peter Railton

INTRODUCTION

My first attempts to drive a car were torture – for myself, my older brother (who unwisely had agreed to help teach me), and the family car. The car had a manual transmission and clutch, and we bucked and lurched around town. Each intersection, even each gear shift, posed a challenge that demanded my full attention – if only I could have given it. Instead, my mind was churning with embarrassment at my incompetence, driven to fever pitch by the chorus of horns that greeted me each time I stalled in traffic. I could barely follow the simplest directions from my brother, and his occasional attempts to calm things down with conversation fell on deaf ears. Despite himself, he groaned quietly as I ground the gears and lugged the engine.

Like everyone, I eventually I got the hang of driving – the way we eventually get the hang of talking, eating without a bib, telling a joke without ruining it, finding our way in a strange city, or politely discouraging an over-eager salesman. What had changed about me as a driver? Not my *rationality*. It was not irrational of me to drive when I was so annoyingly clumsy at it – I had to learn, and there was no other way. My driving was incompetent, but not really dangerous. True, I was responding badly to the available reasons. So I was not a good *detector of or responder to* reasons. But not out of irrationality. Believe me, I was squeezing whatever I could out of reason alone.

What changed was my *competence* or *fluency* as a driver. As I gradually acquired the component skills and gained confidence in my ability,

* I would like to express my appreciation to the editors of this volume, David Sobel and Steven Wall, as well as to those who attended the 2006 Bowling Green conference on Practical Reason, for helpful comments and criticisms. As always, I owe a special debt to my colleagues Elizabeth Anderson, Allan Gibbard, and Stephen Darwall, and to my former colleague David Velleman. Richard Nisbett helped introduce me to recent work in cognitive social psychology.

I became able to drive in traffic without causing bottlenecks and precipitating a chorus of horns. Once skilled enough and sufficiently confident in myself and my skills, I could delegate the task of driving largely to an acquired *habitus*, permitting me to focus on other things while I drove – following complex directions, say, or having an intense conversation with a companion. Skilled driving is not, however, like being on “autopilot” or following a set, mechanized routine. Each drive is different in a variety of ways from the drive before it, and so each presents the driver with somewhat different conditions – somewhat different reasons – to respond to. Moreover, competent drivers do not all drive alike, nor do they drive the same way each time. You can come to know a fair amount about someone or her mood from the way he or she drives. Fluency permits one’s driving to be self-expressive, for better or worse.

Which brings me to the fact that although driving fluency permits better responsiveness to reasons and greater self-expression, it does not guarantee that one’s driving will be rational. If anything, fluency can put one at risk of much more dangerous and irrational driving. When I first began to drive, there was no risk of my threading my way through heavy traffic at excessive speed, or driving when three sheets to the wind.

MISSING PIECES

In philosophy, we find debates between internalists and externalists in virtually every domain, and along a multiplicity of dimensions. Some philosophers, for example, insist that the theory of practical reason should take as its starting point the standpoint of an agent facing a conscious choice. Such *first-personal approaches* require us to focus on whether the agent makes appropriate use of information and motives actually available to her in the situation. This represents one form of “internalism” about practical reasons.¹ Primacy is given to the subjective “ought,” in terms of which any putatively objective practical “ought” will be analyzed. In a competing camp we find philosophers who approach the analysis of rational action from a standpoint “external” to the agent’s own perspective – the *third-personal standpoint* of explanation. For such theorists, the essential

¹ Note that this sort of internalism is distinct from the sort in play in debates over whether motivation is in some sense “internal” to the content of normative judgment (i.e. so-called *judgment internalism*), or a necessary condition for the existence of *bona fide* practical reasons (i.e. so-called *existence internalism*). Here there is a closer analogy with the way the internalism/externalism issue is drawn in epistemology.

thing is to understand the causal structure of intentional action, emphasizing the variety of mechanisms by means of which behavior is responsive to reasons. Primacy is accorded to the objective “ought,” in terms of which any subjective “ought” will be analyzed.

Relatedly, there is a debate about the role of autonomous agency in practical rationality. For some, the very paradigm of practical rationality occurs when an agent deliberates and consciously exerts choice, reviewing upon her reasons, deciding upon their relevance and weight, and making them effective by means of the directed exercise of her will. Contrast approaches that focus in the first instance on the various processes by which our thoughts, feelings, and actions may track, or fail to track, available reasons. For example, individuals engaged in an intense conversation can be observed to mimic one another’s posture, gestures, facial expressions, eye movements, and voice. All this coordination occurs without conscious thought. Normal babies manifest this sort of mimicry from their earliest days, and developmental problems will face those whose show no sign of it. Even in adulthood, such mimicry plays a fundamental role in establishing and sustaining our mutual understanding in social settings. Each of us therefore typically has reason to reciprocate to some degree the movements, gestures, voice, etc. of our conversational partners – failure to do so may impede effective interaction and contribute to misunderstanding. Moreover, failure to notice when the other has stopped mimicking can lead one to fail to see that something has gone awry in the conversation or that one has overstayed one’s welcome. Virtually everyone in society becomes proficient at producing and interpreting “body language,” but vanishingly few are ever aware of doing so. To those philosophers who stress autonomy and deliberate choice in their conception of practical rationality, this sort of acquired, intelligent, reason-responsive conduct is not genuinely rational, because it involves no conscious agency. The reasons to which the individual is responding as seen as *acting on* her, rather than the other way around. To those philosophers who stress the tracking of reasons, however, the notion of a form of agency in which the individual can distance herself from “empirical determination” is suspect on psychological and philosophical grounds – it smacks of a revival of “rational homuncularism” or “agent causation.”

My primary sympathy lies with the second group of philosophers in each pair, above, but there is too much to be said for the other side for me to think that either one gives a full solution to the puzzle of practical reason. Perhaps we philosophers are fighting over how to fit together the

pieces of a jigsaw puzzle when several pieces are missing. Neither side can assemble the puzzle satisfactorily until those pieces are brought to the table. *Agent-competence* and *practical intelligence* are two of those pieces. At any rate, that is what I will argue.

AGENT-COMPETENCE

By “agent-competence” I have in mind something analogous to two more familiar competencies: Linguistic competence and social competence. Despite many differences in the language and culture into which they are born, and despite wide variation in their personal history, virtually all children acquire linguistic competence in a mother tongue. Such competence goes beyond the ability to distinguish grammatical from ungrammatical sentences, and includes a capacity to understand and produce entirely novel sentences, and to use language successfully for meaningful communication. Indeed, by young adulthood, virtually all children have become fluent speakers – able to express complex, conditional thoughts, and capable of producing and understanding sentences at the speed of thought.

That all of this takes place in almost all children and without special training has led linguists to posit a native grammatical capacity, which equips children with a special aptitude for acquiring vocabulary and a tacit knowledge of the syntactic and semantic rules of their language, even though the evidence they will have received from the speech of others is often fragmentary and imperfect. Further evidence of this native capacity comes from the stagewise pattern of linguistic development in infants, and the fact that disease, injury, and birth defects tend to cause discrete, systematic deficits in linguistic abilities, deficits that often are unrelated to general intelligence, and which typically can be overcome, if at all, only with great difficulty. In light of the limited range of formal variation among natural languages and the rapidity with which humans raised in one culture can learn the languages of another, it is supposed that this native capacity has a universal form. The diversity of natural languages arises from separate lexical traditions and grammatical variation along a finite number of parameters.²

² For two views, see M.D. Hauser, N. Chomsky, and W.T. Fitch, “The Faculty of Language: What Is It, Who Has It, and How Did It Evolve?,” *Science*, 298, (2002): 1569–1579 and R. Jackendoff, *Foundations of Language: Brain, Meaning, Grammar, and Evolution* (Oxford: Oxford University Press, 2002).

It is also reasonable to suppose that children have a native capacity for sociability, and, with maturity, for fluent social competence. Newborns typically spontaneously form deep bonds with parents or care-givers, and quickly respond reciprocally to signs of others' pleasure or distress. They track whether others are focusing on them, and learn how to attract and hold attention. These are things that normally need not be, and probably could not be, taught. Perceptual discrimination of intentional vs non-intentional movements and entities appears very early, and without special training. Under ordinary conditions, this social orientation and aptitude leads to development of a sense of self vs others, as well as a greater capacity to understand the thoughts and feelings of others, and to respond appropriately. As in the case of language, we see here a pattern of regular, stagewise, nearly universal acquisition of capacities that together yield a distinctive kind of competence. In most individuals, the end result is a form of fluent sociability in familiar settings. Like language, too, we can observe some global similarities in basic social forms, as well as intercultural intelligibility and individual facility in adapting to the norms and customs of other societies.³ We also find that persistent deficits in sociability tend appear in distinctive patterns, often linked to particular diseases, dysfunctions, injuries, childhood trauma, or birth defects.

Linguistic and social competence are usually called *tacit* competencies since individuals typically lack direct insight into the rules, norms, conventions, etc. they follow as they speak and act. To be sure, individuals do have many "intuitions" – intuitions about the meanings of words, about which sentences are grammatical, about what polite conversation requires, about who is dominant in a given group, or about which behavior is "properly" masculine or feminine. These intuitions, articulate or otherwise, reflect the individual's internalization of local rules, norms, and conventions. The smooth operation of these competencies does not require conscious attention – individuals are guided by an intuitive sense of what is appropriate or anomalous, and of course sheer *habitus* does much of the work in regulating behavior. In acquiring a second language or learning the tacit norms of an organization or society, it is the mark of the transition from plodding competence to genuine fluency that one no longer need mentally "look up" rules or principles in order to speak grammatically or to act *comme il faut*. Fluency is by no means a matter

³ For some sample overviews, see A.P. Fiske, *Structures of Social Life: The Four Elementary Forms of Human Relationships* (New York: Free Press, 1991). Relatedly, see D. Goleman, *Emotional Intelligence* (New York: Bantam, 1995) and *Social Intelligence* (New York: Bantam, 2006).

of sheer “mechanization” or “routinization” alone. Rather, greater fluency typically makes speech or behavior less mechanical, more personally expressive, and more open to innovation. Indeed, highly fluent individuals have a good sense of when it is appropriate to bend or break the rules.

Linguistic and social competence, and, more generally, the capacity to acquire the values and ways of life of the community into which one is born, are so central to human life that we can expect natural selection to have equipped us with the wherewithal for the reliable acquisition and fluent deployment of these interlocking skills. Each competence is, moreover, sufficiently integrated functionally to warrant speaking of them as *psychological natural kinds*. (Some psychologists speak of these competencies as involving “modules,” though that is a much more controversial claim.)

My speculative hypothesis is that the capacity to *act intentionally* and *interpret the actions of others* is likewise an essential competency. Let me dub the cluster of abilities that creates the ability to exercise effective intentional agency and interpretation *practical competence*. This is, I believe, like *intentional action* itself, a psychological natural kind, functionally integrated and bound to have been evolutionarily shaped and favored.

The first component of practical competence is what I will call *agent-competence*. Roughly, this is a set of skills akin to the skills that equip someone to be an effective administrator, *inter alia*: an ability to focus attention selectively, but also to be somewhat mindful of several things at once; an ability to set goals and make plans, but also to revisit and revise them in light of experience; an ability to interpret situations and form a view about what’s at stake and which actions are in the offing; an ability vividly to imagine alternatives and identify necessary means toward them as well as their remoter consequences; an ability to deliberate and decide among options; an ability to adjust means to ends; an ability to maintain morale and generate motivation in the face of difficulty; and so on.

Many of the elements in this list might be thought to be part of practical rationality. How is agent-competence different from practical reason itself? Practical reason, we might say, represents a distinctive way of *deploying* of one’s agent-competencies, just as rational driving is a distinctive way of deploying one’s driving competence and rational discourse and social conduct are distinctive ways of deploying one’s linguistic and social competences. Developmentally, agent-competencies emerge over time, and are important *precursors* to, and *conditions* for, mature practical rationality. However, although they play an important role in

enabling the exercise of effective rationality, their successful development, even to the point of fluency, is no guarantee of rational conduct. Highly effective administrators have, for example, been keenly involved in carrying out some of the most senseless or horrible economic, military, and political strategies.

Garden-variety irrational actions – those involving from weakness of the will, failing to consider alternatives or remoter consequences, indecision, or irresoluteness, failure to adjust means to ends, etc. – often involve some measure of failure in agent-competence. But even in our irrational moments, we typically retain a sufficiently high level of agent-competence to give our actions some efficacy. A warehouse employee with a prior record of arrest for shop-lifting understands that he *must not* pilfer – management is keeping an eye on him, and a second arrest on his record would ruin his hopes for a decent job. He knows that he can't blow this chance. Yet when a shipment of sleek, new iPods comes in, temptation sets in. While stocking the shelves, he stacks cartons high enough to block the hidden camera as he bends down to open a fresh box of iPods, and deftly palms one and slides it under his shirt. His theft, if successful, will net him one rather expensive and useful gadget; but if detected, it will cost him his job and much more. Were he fully rational, he wouldn't risk it. But he isn't fully rational – temptation gets the better of him. Still, his irrational theft is enacted very competently. He does not attempt to steal something of very low value or impossible to conceal, and he accurately envisages how he might be detected and acts accordingly. He goes through with it with the required feigned nonchalance, and continues about his task as if nothing had happened. All of this is his agent-competence at work – even in the name of an aim a rational person in his shoes would have rejected outright. Sheer competency cannot save us from ourselves.

We might contrast this employee's behavior with that of a young child stealing from the change bowl. After spotting the change bowl on a low shelf, he may approach it hesitantly, conspicuously looking around to see if he is being watched, pausing with his hand poised over the change bowl still trying to make up his mind and momentarily stymied by the unanticipated problem of deciding which coins to take. He then grabs a handful of coins noisily, still looking around nervously, and darts out of the room in a state of high nervous agitation and with a very guilty look on his face. A parent or friend encountering him in the next room would sense at once that something was up. His pilfering would be as incompetent as my early driving.

At bottom, the child and the adult succumbed to remarkably similar psychic forces. Each knew better than to steal in the circumstances, and each was swayed by the presence of a highly salient stimulus. In that sense, the adult behaved “like a child.” Yet in another sense, the adult’s behavior was not childish at all, and deserves higher marks for execution. We would, in effect, see a *double* lack of mature agency in an adult who, in attempting to pilfer the iPod, went about it just like the young child.

It might seem that we prefer personalities, even among adults, inimical to skillful wrongdoing. Good people, one might think, will by nature make bad thieves. After all, accomplished thieves seem disturbingly cold-blooded and indifferent to the harm they cause others, especially to those most vulnerable individuals who often are their targets. Perhaps skill at theft deserves nothing but contempt? This pious thought is not sustainable. There will be cases in which we need to steal for a good cause – to filch food from the rations stored in a municipal warehouse in wartime in order to provision the Resistance, say. In such circumstances, we may especially *admire* someone who, even though brought up never to steal and hating the thought of putting her fellow employees at risk, nonetheless has the self-mastery, *sang froid*, and shrewdness to carry out the theft with the fluency needed for success. This is no time to be weak-kneed, or paralyzed by conventionality. Yet admirable as such self-mastery, *sang froid*, and deftness may be, I doubt we see them as part of practical rationality. We do not, for example, lower our opinion of the practical rationality of otherwise estimable agents who simply could not pull off this off.

Agent-competence lends structure and meaning to behavior, permitting us to see the intentions of the agent behind it – even when the intentions are irrational or malevolent. Trying to understand the concern behind the flailing and wailing of an overwrought child is difficult precisely because the child is failing to manifest agent-competency – rather like the case of trying to interpret the speech of a highly agitated individual lacking basic competence in our language. Figuring out what is bothering an overwrought child often requires finding some way to bring into play the competencies he *does* possess, even when distraught. “Can you *point* to what’s bothering you?” the nurse asks in a reassuring tone.

PRACTICAL INTELLIGENCE

The second component of practical competence, and the second piece of the puzzle of practical rationality that I would like to bring to the table, is

what I will call *practical intelligence*. One credible way of characterizing “general intelligence” is *problem-solving ability*, and we can think of practical intelligence as the ability to solve a range of problems specific to the exercise of agency. Practical intelligence will often be needed in order to possess or exercise practical competencies.

A paradigmatic example is possession of the ability to contend with the problem of impulsivity in the face of temptation, as in the example above. What equips us for it? In a classic series of experiences, Walter Mischel tested the ability of four-year-olds to resist eating a marshmallow placed before them, on the strength of a promise that if they could refrain from eating it for the next quarter-hour, they would receive a second marshmallow. This sounds like a job for practical reason and “will power,” but the four-year-olds who resisted impulse successfully relied less on iron logic and steely resolve than on indirect strategies of distraction or disengagement – pretending the marshmallow was something else, placing their head down on the table to hide their eyes, concentrating on other things, trying to fall asleep. The central importance of this capacity to find a way of *working around* or *working with* impulse – since it is a motivational state one cannot cancel at will or refute by argument – became manifest when the same individuals were studied again in grade school and then, still later, in high school. Those unable to work around or with their temptation, who ate the first marshmallow soon after the test began, were found to have had more problems with social adjustment, behavior control, peer relationships, and schoolwork.⁴

Importantly, the skill in question is not sheer “self-denial” or “self-mastery,” though it is a skill that may facilitate denying oneself and mastering one’s feelings when this serves a purpose. Just as those who immediately indulged themselves had greater subsequent problems in life, so did those who stoically soldiered on in self-denial, resisting the second marshmallow once it was produced and failing to reward themselves by eating the marshmallows, once there was nothing to gain from further delay.

Moreover, the skill involved does not appear to be excellence in practical reasoning. True, those who managed to fend off impulse, and to enjoy the reward of success once achieved, were more responsive to available reasons than the others. But the explanation does not appear

⁴ Y. Shoda, W. Mischel, and P.K. Peake, “Predicting Adolescent Cognitive and Self-Regulatory Competencies from Preschool Delay of Gratification,” *Developmental Psychology*, 26, (1990): 978–986.

to lie in their reasoning ability or recognition of what was required or prudent in the situation. All the children liked marshmallows, and all preferred two over one. All understood the point of the game and the advantages of waiting. In that sense, the children all saw the same reasons to consume the first marshmallow and the same, stronger, reasons to delay. The differences among the groups appear to have arisen from a bundle of non-deliberative cognitive, imaginative, and affective abilities. It was possession of this bundle, it appears, that proved so valuable in so many areas of their lives as they grew older and faced new problems.

Research on moral development also suggests that reasoning capacity and grasp of principles may be less central than a capacity for empathy and emotional engagement. Children's behavior in actual choices, for example, was better predicted by these affective capacities than by the complexity of their moral reasoning or ability to enunciate principles.⁵ Similarly, among a group of Polish policemen ordered by Nazi occupiers to execute civilians at point-blank range, those who succeeded in non-compliance were those who seem empathetically to have mirrored the fear and panic of their victims and experienced a visceral revulsion at shooting. By contrast, few cited moral principles or religious convictions to explain their conduct.⁶

These characteristics – empathy and visceral resistance – are in some sense pre-moral or proto-moral, rather than “moralized” or “judgmental.” Empathy, for example, can operate contra-judgmentally. It will induce alarm and distress in response to the terror seen in others' faces even when those exhibiting terror are judged to be one's enemies or to deserve punishment. Similar characteristics, and similar neurological structures – a similar chemistry of “bonding,” “cooperativeness,” and “inhibition toward harming con-specifics who signal helplessness” – can be found in social

⁵ See J. Metcalf and W. Mischel, “A Hot/Cool System Analysis of the Delay of Gratification,” *Psychological Review*, 106, (1999): 3–19; M.L. Hoffman, “Development of Prosocial Motivation: Empathy and Guilt,” in N. Eisenberg, ed., *The Development of Prosocial Behavior* (New York: Academic Press, 1982); J. Haidt, “The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment,” *Psychological Review*, 108, (2001): 814–834; A. Damasio, *Descartes' Error* (New York: Putnam, 1994); and C.D. Batson et al., “Value Judgments: Testing the Somatic-Marker Hypothesis Using False Physiological Feedback,” *Personality and Social Psychology Bulletin*, 25, (1999): 1021–1032. One interesting qualification: Moral reasoning ability was more strongly linked with successful inhibition of delinquent behavior than successful initiation of helping behavior. See S.J. Thoma et al., “Does Moral Judgment Development Reduce to Political Attitudes or Verbal Ability?,” *Educational Psychology Review*, 11, (1999): 325–341.

⁶ For discussion, see R.F. Baumeister, *Evil: Inside Human Violence and Cruelty* (New York: W.H. Freeman, 1997): 209–210. Baumeister's analysis is based upon the research of Christopher Browning, *Ordinary Men: Reserve Police Battalion 101 and the Final Solution in Poland* (New York: HarperCollins, 1992).

animals, who presumably lack a capacity for practical reason. So, too, do we find in social animals and humans alike “mirror neurons” that replicate the behavior of others, and brain systems for the empathic emotional “simulation” of the states of others on a substrate of the individual’s own affective system. Humans may differ from animals less in their capacity for such responses and feelings, than in their ability to mobilize or (regrettably often) *de*-mobilize these responses in the name of abstract ideals or conceptions of “sameness” vs “otherness,” or “superiority” vs “inferiority” – and, of course, in the magnitude of the harm they can do when some of these systems are absent or deficient.⁷

The affective and emotive systems that appear to play a pivotal role in our thought and conduct are informed about the world directly from the senses, not only through conscious experience, but also via dedicated, fast pathways unmediated by higher-order cognitive interpretation. Typically, they come into play even before cognitive judgments are formed, and they characteristically prime and color such judgment. Individuals who suffer from impairments in these affective and emotive systems, or in the pathways by which they send signals to higher-order cognition, appear to have greater difficulty forming accurate judgments of the risks they face, and to show lesser ability to track their own risk judgments effectively in the choices they make. Moreover, they more likely to experience difficulties in social relations, or to exhibit asocial or sociopathic behavior.⁸

This *primacy of affect* is a fundamental and pervasive feature of our psyche. It probably is no exaggeration to say that the bulk of our everyday mental activity – perception, attention, memory, learning, inference, association, intention-formation – is affect-mediated.⁹ The role of affect in shaping subsequent judgment and behavior is sometimes called the “automatic evaluation process,” and the picture of the psyche that emerges is known as a “dual-process” model, with fast affective processes and slower cognitive processes running simultaneously in the mind.¹⁰ Information

⁷ Practical reason *can* play an important role here, for example, by alerting us to these risks, and by challenging conceptions of “otherness” or otherwise encouraging the development of empathic responses. In this way, humans can be “argued into” sending aid to the other side of the globe, or passing and enforcing laws to ban violence or discrimination. No animals have yet shown these abilities. At the same time, no animals have ever shown violence toward con-specifics on the human scale.

⁸ See Damasio, *Descartes’ Error*, esp. chapter 3. The implications of this for “judgment internalism” are briefly discussed below.

⁹ The seminal work is R. Zajonc, “Feeling and Thinking: Preferences Need No Inference,” *American Psychologist*, 35, (1980): 151–175. For more discussion, see Haidt “The Emotional Dog.”

¹⁰ For discussion, see J.A. Bargh and T.L. Chartrand, “The Unbearable Automaticity of Being,” *American Psychologist*, 54 (1999): 462–479.

arriving directly from the senses is compared with expectations and immediately coded in the brain as favorable, neutral, or unfavorable. Because the coding is affective, it not only frames judgment, but also sets in motion motivation and emotion.¹¹ Thus a judgment and a corresponding emotive or motivational response will tend to co-occur, not because either one constitutes the other, but because they are, in part, effects of a common cause. The first stages of these processes occur with remarkable speed, at a level unavailable to conscious introspection. From the standpoint of our self-awareness, the flow from perception to intuitive judgment is “seamless” – we see a dangerous condition *as* dangerous, immediately feeling the peril it poses and preparing for action.

Moral judgment may be an example of just this phenomenon. The psychologist Jonathan Haidt has studied human moral judgment by presenting subjects with hypothetical scenarios and recording their reactions. He finds that subjects presented with a moral scenario involving incest, for example, have a quick affective response, pro or con, which precedes and appears to prime conscious judgment. Because the origin and occurrence of this intuitive affective response is not introspectively available, subjects asked to explain *why* they judged the scenario as they did must improvise a plausible rationale. Of course, they do not sense that they are constructing a rationale *post hoc* – it seems to them that they are giving their actual reasons, and they do tend to cite considerations of a kind that *would*, if present in the scenario, explain their judgment. However, the scenario has been carefully designed so that most of the considerations it initially occurs to people to cite are not in fact present. Thus, the scenario stipulates that the incestuous act took place between mature adults who cared for each other, acting fully voluntarily and without coercion, avoiding the possibility of conception, and suffering no subsequent psychological trauma. When the disparity between the scenario actually presented and the rationale they provide for their judgment is pointed out to subjects, they tend *not* to withdraw or reconsider their initial judgment.¹² This firmness of judgment in the absence of a rationale would make sense if the explicit judgment and the *post hoc* rationalization have a common cause, namely, the intuitive affective

¹¹ Fiske introduced the term “category-based affect” to describe this phenomenon. See S.T. Fiske, “Schema-Triggered Affect,” in M.S. Clark and S.T. Fiske, eds., *Affect and Cognition* (San Diego: Academic Press, 1982).

¹² Haidt, “The Emotional Dog.”

response. Since the intuitive response never depended causally upon the rationale, it is unaffected by its collapse.

The primacy of affect over reasoning in giving rise to judgment makes sense.¹³ We are after all the products of millions of years in which progenitors less well endowed with higher-order cognitive resources were selected for their ability to rapidly and reliably discriminate the propitious or beneficial from the risky or harmful. Studies of primate neurology support the hypothesis that the brain makes a fast “triage” of incoming information, as positive, neutral, or negative, priming the animal’s responses accordingly. At the same time, the brain updates prior expectations according as the new information is better than, worse than, or the same as expected. In this way the “triage” process is constantly *learning* rather than operating with a fixed standard. In effect, it learns from experience how to do its triage more accurately and guide the organism toward its goals more reliably.¹⁴ Humans appear to have retained this “automatic evaluation system,” the power of which is increased by the range and complexity of the expectations we can form and the goals we can pursue. Those with highly developed practical skills – in sports, seamanship, teaching, etc. – appear to draw upon this evaluation system, as refined by experience, in forming a quick intuitive “sense” of what is happening in a given situation, and how to respond appropriately.

AFFECTIVE REGULATION:
IT’S NOT JUST FOR HUMEANS ANY MORE

The pervasive regulatory role of affect in our mental life is, of course, a favorite theme of Humeans, but the phenomenon did not altogether escape the founders of other traditions in the theory of practical reason. Kant located the will in the “faculty of desire,” rather than some separate, higher faculty.¹⁵ Moreover, he characterized the empirical psychology of a virtuous agent who respects the moral law in affective terms – our respect for the moral law is “the moral feeling,” not a pure cognition.¹⁶ Our sense of the

¹³ Haidt found an exception to this pattern in the case of moral judgments when subjects are presented with a classic “moral dilemma,” which makes reasoning salient. This phenomenon is discussed further, below. See Haidt, “The Emotional Dog” and S. Murphy, J. Haidt, and F. Bjorklund, “Moral Dumbfounding: When Intuition Finds No Reason” (pre-print).

¹⁴ See W. Shultz, P. Dayan, and R.P. Montague, “A Neural Substrate of Prediction and Reward,” *Science*, 275, (1997): 1593–1599.

¹⁵ I. Kant, *The Metaphysics of Morals*: trans. M.J. Gregor (Cambridge: Cambridge University Press, 1966): 6:213.

¹⁶ *The Metaphysics of Morals*: 6:402.

beautiful or sublime in both morality and nature is immediate, and “not brought to concepts,” even though it guides our respect, strikes down our self-conceit, and produces our appreciation and awe.¹⁷

Aristotle, for his part, stressed that virtue is not a matter of conscientiousness and rational application of principle alone – for such is the nature of mere continence – but of right feeling and finding happiness in the right things. Affect, he thought, had the fundamental regulatory role in our psyche:

[A]nd to a greater or lesser extent, we regulate our actions by pleasure and pain. Our whole inquiry, then, must be concerned with them, because whether we feel enjoyment and pain in a good or bad way has great influence on our actions.¹⁸

The phenomenon of regulation by affect (one example of which is Aristotle’s treatment of pleasure and pain as “standards”) is exemplified at its most basic by classical conditioning. Teaching animals – whether rats in a cage or elephants at the circus – invariably involves the carefully regulated use of reward and punishment, in effect redeploing the animal’s internal affect-governed regulation process on behalf of an agenda set externally by the trainer.

This mention of the training of animals will surely prompt thoughts that an account of psychology assigning a central role to affect regulation “reduces us to mere animals.” But this would be “reducing” only if the things in which humans find intrinsic reward were animalistic. They are not. Humans find reward in pleasant experience, to be sure, but also in such things as successful pursuit of abstract ideals, excellence in the exercise of skills and capacities, discovery and the creation of knowledge, friendship, humor, self-expression, aesthetic appreciation, romantic love, and commitment to kith and kin. Human affect regulation can be as diverse and conceptually rich as our desires and passions themselves.

PRACTICAL COMPETENCE

Human affect regulation is part of what some psychologists have come to call our “emotional intelligence.” This intelligence includes a sense of one’s own feelings and wants as they occur, an ability to sense and anticipate the thoughts and feelings of others, an ability to motivate oneself on behalf of a goal, an ability to control impulse and modulate

¹⁷ I. Kant, *The Critique of Judgment*, trans. W.S. Pluhar (Indianapolis, IN: Hackett, 1987): 266.

¹⁸ Aristotle, *Nicomachean Ethics*, trans. R. Crisp (Cambridge: Cambridge University Press, 2000): 1105a3–5.

or channel one's emotional responses, and an ability to form bonds with others and manage emotion within a relationship.¹⁹ Evidence ranging from health indicators to measures of financial success and personal achievement and life-satisfaction supports the view that this sort of intelligence may be more important for one's well-being, effectiveness, and overall life prospects than cognitive intelligence.²⁰ It is readily seen that this sort of intelligence would be important part of our intelligence *as agents*—that is, of our practical intelligence.

Practical intelligence is largely a matter of *know-how* rather than *knowledge-that* or reasoning ability – knowing how to hold off impulse, how to sustain one's motivation and protect one's morale, how to plan, how to play one's part in a conversation, how to reciprocate affection, how to acknowledge and learn from one's mistakes, how to apply the information one has, and so on. Moreover, it often involves something like “opponent processing.” Alongside knowing how to plan and stick to a plan is knowing how to avoid crippling rigidity by remaining open to challenge, change, unanticipated experience, and emergent opportunities. Alongside knowing how to hold off impulse is knowing how to escape untempered asceticism, rewarding oneself effectively and enjoying life's pleasures. Alongside knowing how to face up to one's mistakes and limitations is knowing how to move on and not be too hard on oneself. Most such know-how cannot be codified as rules. It is acquired, if at all, through experience and emulation, and it manifests itself in our lives through the fluent and effective exercise of competencies.

Taking together agent-competence and the practical intelligence, then, we arrive at what might be called *practical competence*.

PRACTICAL COMPETENCE AND AUTOMATICITY

Fluencies such as skilled driving, speaking one's native tongue, or adapting rapidly to the mood or manners of a group are often called “automatisms” by psychologists, but this term may misleadingly suggest that the operation of these fluencies is mechanical, inflexible, or unfeeling. In truth the exercise of these fluencies is, or at any rate can be, extraordinarily well attuned to circumstance and mood. Once we have become fluent, walking upright and speaking a second language are “automatic”

¹⁹ See P. Salovey and J.D. Mayer, “Emotional Intelligence,” *Imagination, Cognition, and Personality*, 9, (1990): 185–211.

²⁰ For an overview, see Goleman, *Emotional Intelligence*.

behaviors in the psychologists' sense, but we do not walk and talk like robots following set programs. Instead, it is our very fluency that permits these acts to be expressive, even creative.

"Automatized" action goes far beyond the *reflexive* or *instinctual*. Reflexes – eye blinks, knee jerks, etc. – involve dedicated circuitry that produces fixed behaviors, often without much help from the central nervous system. Instincts – fear of spiders and snakes, mating rituals, homing behavior – typically manifest themselves in stereotyped behavior elicited by a fixed array of stimuli. Reflexes and instincts are also typically resistant to modification through learning. In contrast, most of the fluencies we have been discussing are intelligent and plastic – they involve extensive learning, anticipating novel problems, generating novel solutions, and a capacity for self-revision. For example, "automatic" eye movements acquired through experience and development can be faster and more subtle learners than our "higher" cognitive faculties, often instructing or guiding our more articulate cognition in quite sophisticated ways. In the "false belief task," a child subject is told a story about two girls, Sally and Anne. Sally has placed a piece of candy in a basket and left the room. Anne, in Sally's absence, removes Sally's candy from the basket and places it in a box. The child subject is then asked where Sally will look for her candy upon returning to the room – the basket or the box? Success at this task requires the ability to impute to another a belief about the world that one does not oneself possess, and to use this imputed false belief to predict and explain the other's behavior. Normal children usually do not succeed at this task until age four. But even before children still have mastered the verbal task, while they are still answering, "In the box," their eyes have begun to dart momentarily toward the basket prior giving their response. This eye movement is nothing the child is aware of, and certainly not voluntary in the usual sense. Yet the eye movement does appear to cue the later emergence of declarative knowledge that Sally will look in the basket, and is absent in children with forms of autism that impede mastery of the false belief task. This absence in turn is correlated with long-term deficits in their ability to understand the feelings and behaviors of others, deficits that even high intelligence does not seem to remedy fully.²¹

²¹ See S. Baron-Cohen, *Mindblindness: An Essay on Autism and Theory of Mind* (Cambridge, MA: MIT Press, 1995). "Solving" the false belief task described above of course involves many other skills on the part of the child, including notably the ability to handle complex verbal instructions involving embedded propositional attitude ascriptions. Non-verbal forms of the task have been developed, and some evidence suggests that children understand them at an earlier age. See P. Mitchell, *Introduction to the Theory of Mind: Children, Autism, and Apes* (London: Arnold, 1997).

The amount, variety, and detail of information gathered by our “intuitive” processes often exceed anything conscious deliberation could manage.²² In one recent study, although subjects facing a hypothetical choice with a small number of variables did better (according to their own announced *ex ante* preferences) when they were given peace to concentrate and time to deliberate. But when the number of variables grew, the same individuals made better choices (again, by their own lights) if they were distracted and allowed insufficient time to deliberate, so that they were forced to rely more heavily on intuition or hunch.²³

The psychologist Jonathan Haidt writes:

It is now widely accepted in social and cognitive psychology that two processing systems are often at work when a person makes judgments or solves problems. Because these two systems typically run in parallel and are capable of reaching differing conclusions, these models are usually called *dual-process* models.²⁴

One of these systems, Haidt notes, is fast, intuitive, effortless, and its operation is typically not available to introspection. The other is consciously aware, slower, more deliberate, and effortful. Each has its strengths and weaknesses, and their operation in tandem is a distinctive feature of our human ability to respond to reasons in thought and action.

As a result, “automatized” behavior can be found across the entire span of human activity. The resulting actions, while not anticipated by conscious intention, may nonetheless be intentional, and done for reasons. The jazz saxophonist’s solo riff, the basketball guard’s well-timed jump, the experienced driver’s smooth downshift, and the wit’s lightning riposte aren’t *unintentional* or *mindless* behavior, like absent-mindedly tapping one’s foot while writing or succumbing without realizing it to the emotional contagion of a crowd. Rather, they are complex, structured, purposeful activities done mindfully but fluently, without deliberation or intention-formation. Yet were we foolish enough to interrupt these individuals in mid-stream, they could typically answer the Anscombe question, “What are you doing?,” without any further observation or inference.

The saxophonist’s improvisations, the guard’s jumping, the timely downshift, and wicked comment are clearly done *for reasons*, and,

²² For examples, see R.R. Hassin, J.S. Uleman, and J.A. Bargh, *The New Unconscious* (Oxford: Oxford University Press, 2005).

²³ See A. Dijksterhuis *et al.*, “On Making the Right Choice: The Deliberation-without-Attention Effect,” *Science*, 311, (2006): 1005–1007.

²⁴ Haidt, “The Emotional Dog”: 819.

moreover, for reasons *as such* (rather than, say, through a deviant causal path). The saxophonist suddenly drops a register from an active, intuitive sense of how the music and mood are developing, the guard jumps how and when he does in order to block a hook shot, the driver shifts down to get more power in a tight curve, and the wit's words fit the occasion only too well. Far from being merely swept along by a causal chain or robotically enacting a habit or routine, the jazz improviser and skilled basketball player, for example, are exercising agency in a pure form, fully deploying their cognitive and creative skills.²⁵ Of course, these aren't examples *par excellence* of deliberative agency. But why see deliberative agency as the truest form of acting for the sake of reasons? Why not see deliberative agency as one more domain – distinctive and important, but by no means predominant – in which humans can develop greater or lesser skill at responding aptly to reasons? After all, deliberation itself is embedded in a host of reasons-responsive activities on the part of the agent, and can only occur thanks to them. Indeed, even being adept at when to deliberate is just such a non-deliberative skill.

AUTONOMY AND AUTOMATICITY

But wait. For all its merits as a way of getting things done well in response to reasons, isn't "automatized" action nonetheless not autonomous? After all, by its nature its operation is not under direct management by the conscious self. The jazz improviser, for example, may be at a loss to explain just how it occurred to him to shift suddenly down a register or to slow down the tempo; the guard who has just blocked a shot may be unable to say just how he "decided upon" which player would receive the pass, or why he leapt at just the instant he did, or how to raise his arm and spread his fingers in order that they meet the ball. Skills fluently exercised are certainly not be robotic, but they do seem to "have a mind of their own," one that does not require intervention by an agent's deliberation, decision, volition, or even approval. The spontaneous working of a skill in action seems in this way to possess a certain self-sufficiency or auto-nomy *within* the agent, thereby compromising *his* autonomy as an agent.

A compelling case can be made that the purposive operation of one's unself-conscious self is not the exercise of one's own agency. Passing

²⁵ For a description of such phenomena, and their ties with optimal experience and self-efficacy, see M. Csikszentmihalyi, *Flow: The Psychology of Optimal Experience* (New York: Harper and Row, 1990).

beyond examples like exercising a skill, David Velleman gives the following example, which is worth quoting at length:

Suppose that I have a long-anticipated meeting with an old friend for the purpose of resolving some minor difference; but that as we talk, his offhand comments provoke me to raise my voice in progressively sharper replies, until we part in anger. Later reflection leads me to realize that accumulated grievances had crystallized in my mind, during the weeks before our meeting, into a resolution to sever our friendship over the matter at hand, and that this resolution is what gave the hurtful edge to my remarks. In short, I may conclude that desires of mine caused a decision, which in turn caused the corresponding behavior; and I may acknowledge that these mental states were thereby exerting their normal motivational force, unabated by any strange perturbation or compulsion. But do I necessarily think that I made the decision or that I executed it? Surely, I can believe that the decision, though genuinely motivated by my desires, was thereby induced in me but not formed by me; and I can believe that it was genuinely executed in my behavior but executed, again, without my help. Indeed, viewing the decision as directly motivated by my desires, and my behavior as directly governed by the decision, is precisely what leads to the thought that as my words became more shrill, it was my resentment speaking, not me.²⁶

Velleman is careful to add that he is not trying to escape responsibility for the blow-up, since he takes himself to be obliged to be “vigilant against unconsidered intentions” and show self-restraint when angered. However, there will be circumstances in which such vigilance and self-restraint are not a practical possibility. We can imagine that things “happened too quickly” when he began talking with his friend. Involuntarily offended by his friend’s manner, a wounding remark escapes his lips before he realizes what he is saying. His friend immediately retaliates with a yet more hurtful remark. Escalation and eventual rupture are now inevitable. Despite their good intentions, the meeting appears to have been hijacked by their feelings. Velleman writes, if his behavior:

could come about *only* in the manner described here – that is, springing directly from intentions that have simply come over me – nothing would owe its occurrence to either my participating or failing to participate in events, and I might bear no responsibility for anything.²⁷

We do indeed have an intuition that supports this notion of the self, and that makes his suggestion that “it was my resentment speaking, not me”

²⁶ J.D. Velleman, “What Happens When Someone Acts?,” reprinted in *The Possibility of Practical Reason* (Oxford: Oxford University Press: 2000) 126–127.

²⁷ Velleman, “What Happens When Someone Acts?”: 127n., emphasis added.

plausible. However, I'd like to suggest that this intuition should be taken with a grain of salt.

Let's begin with an account of intentional action that lays out an explicit role for agency. In *Rationality in Action*, John Searle claims:

Consider any situation of rational decision making and acting and you will see that you have a sense of possibilities open to you . . .²⁸

He contrasts such situations with those in which, for example, you are in the grip of anger and have "no sense that you could be doing something else." There is no distance between how you feel and what you do. You have "lost control" and become a channel for your anger. A similar sense that what happens is no longer "up to you" can be experienced by those in the grip of a powerful addiction or compulsion.²⁹

So *bona fide* intentional agency combines a sense of open possibilities with the fact that *which* possibility is realized will be determined by your choice. On this account, we can see an essential role for the self, since choice must insert itself wherever there are "gaps" in agency, and Searle mentions three such gaps.

First, there is a gap between the reasons you consider and the decision you actually make – even when the decision is based upon those reasons. Once you have satisfied yourself that you have identified a "rationally sufficient reason," for example, your job as rational agent is not done. This sufficient reason will not enact itself: you must make up your mind and do the deciding about what will be done and why. In so doing, you make it *your reason* – you choose the reason for which you act, good or bad.

Second, there is a gap between your decision and the action itself, since (except in some special cases) decisions in themselves do not ordinarily causally suffice for initiating the action decided upon. Having decided, you typically cannot sit back, put up your feet, and wait for the action to occur. Here, too, as Searle writes, "you actually have to do it." In so doing, you make this *your doing* – the act that you yourself have undertaken, wisely or foolishly.

Third, there is a gap between beginning an action in this way and carrying it through to completion. Since many actions extend over time, this gap is familiar to all of us. Getting yourself off the couch and doing your nightly exercises may be a real feat, but once underway you cannot simply rest on your laurels. You cannot say, "Well, the die is cast – my

²⁸ J. Searle, *Rationality in Action* (Cambridge, MA: MIT Press, 2001): 15.

²⁹ Searle, *Rationality in Action*: 15–16.

nightly exercises are underway.” Exercising is not like jumping from a building – one cannot start the action and then simply watch events unfold as the causes operate by themselves. The exercising stops the instant you do. Filling this gap – really, these gaps – calls upon your agency once more. Even as your exercises tire you, “you have to make a continuous voluntary effort to keep going with the action.”³⁰ In so doing, you make this *your accomplishment*, for good or ill.

Practical rationality, Searle claims, operates in these gaps. And when it does, it cannot look around for *something* to plug the gap. In that sense, he writes, “Nothing fills the gap.” Rather:

you make up your mind to do something, or you just haul off and do what you are going to do, or you carry out the decision you previously made, or you keep going or fail to keep going . . .³¹

Let us take an example Searle offers of action for a reason:

Suppose you are asked to justify voting for Clinton; you might do so by appealing to his superior management of the economy. But it may be the case that the actual reason you acted on was that he went to your old college at Oxford, and you thought, “College loyalty comes first.” And the remarkable thing about this phenomenon is: in the normal case you know without observation which reason was effective, because you made it effective. That is to say, a reason for action is an effective reason only if you made it effective.³²

We have already seen what making a reason *your* reason involves:

The reasons did not operate on you. Rather you *chose* one reason. You made that reason effective by *acting on it*.³³

By contrast, in Velleman’s example involving loss of temper, we might say that the agent did not choose the reasons that led to the irruption, did not choose to start a fight, and did not choose to carry the quarrel to its rupturing conclusion – the reasons for the quarrel and rupture “operated on” the agent, not he on them. The typical case of intentional action, Searle writes, is quite unlike this. Instead: “Deliberation typically leads to intentional action by way of prior intentions.”³⁴

To be sure, Searle recognizes that not all intentional activity follows this canonical model of deliberation leading to a “choice of one’s reasons”

³⁰ Searle, *Rationality in Action*: 14–15.

³¹ Searle, *Rationality in Action*: 17.

³² Searle, *Rationality in Action*: 16.

³³ Searle, *Rationality in Action*: 16.

³⁴ Searle, *Rationality in Action*: 47.

and to formation of an intention. Some intentions are not formed consciously, like Velleman's intention to break off the friendship, and "not all" intentional action is premeditated – some is "spontaneous," like rubbing one's head or rising to one's feet to pace when struggling with the wording of an argument.³⁵

What I want to argue is that, as far as I can see, far from being typical of intentional action, Searle's model does not fit the vast majority of everyday intentional actions. Such actions are not deliberated or premeditated, but not for that unintentional. Haidt writes:

The emerging view in social cognition is that *most* of our behaviors and judgments are in fact made automatically (i.e. without intention, effort, or awareness of process).³⁶

Whether Haidt is right or wrong is an empirical question, one about which I have no special expertise. Instead, I wish to argue something stronger, namely, that *all* action – including in particular paradigmatic premeditated intentional action – has *and must have* unpremeditated action at its source and core. A corollary: Most of the reasons for which we act, and that give us the name of rational beings, are not made effective by "choosing one's reasons."

One way to see the problem is to consider again what Searle says about "filling gaps" in the enactment of an intention. What he proposes to fill these gaps in action is each itself an action, a *doing* – my "actually deciding," or my "hauling off and doing," or my "making a continuing voluntarily effort" – and, presumably, a *doing for some reason*. But now, if these are to be full-fledged intentional actions of mine, attributable to my agency in virtue of conforming to Searle's model, each must itself have the same three gaps to be filled.

For example, consider the first gap in Searle's example of facing a choice between Clinton and Bush. The first gap to be filled, Searle explains, is that the agent must "choose his reason" for deciding – the reason cannot simply "act upon him." It must truly be *his* reason, made effective by him. In Searle's example, the fact of sharing an *alma mater* with Clinton is the reason the agent chooses to decide for whom to vote. Very well, but if this is to be a *choice* on the part of the agent, not an arbitrary picking or a reason "acting upon him," this selection of alumni

³⁵ Searle, *Rationality in Action*: 45.

³⁶ Haidt, "The Emotional Dog": 819.

status (as opposed to myriad other potential reasons) as the ground for his choice must itself be done by the agent for a reason of his own choosing. That is, on Searle's scheme, it would seem that the agent's choice of Clinton could only be autonomous – *his* choice, for *his* reasons – if he were to choose this reason, too. But now we are on our way to a regress – instead of one gap to fill, we have two. And so on.

Of course, the problem faced here isn't peculiar to Searle's three-gap model. It arises for any model of action that seeks to understand the distinctive operation of autonomous or rational agency in terms of some special sort of action on the part of the agent, whether the act is "choosing one's reasons," or "endorsing certain reasons," or "identifying a certain reason," or "throwing one's weight behind one reason rather than another." Since it would appear that these acts would themselves have to be done autonomously, the would-be agent has become *Zeno's deliberator*.

DELIBERATING WITHOUT REGRESS

In practice, of course, we do somehow manage to deliberate and act intentionally without losing ourselves in regress. How is this possible, consistent with the idea that the act we decide upon, and its reasons, really are our own? We have seen how certain initially plausible "top-down" solutions would generate a regress unless there were some way for an agent to choose, or endorse, or identify with, or throw his weight behind certain reasons autonomously, but without an act of deliberate choice. Suppose there were, in other words, ways in which individuals could come to embrace one reason over others autonomously, but not via a further "full fledged" act? Then regress might be avoided. But then, too, the privileged status of deliberate choice as a model for autonomy would have been undermined, and we would have to ask, "If in these cases an act can be autonomous without involving deliberate choice, why not in other cases?"

I have no quarrel with treating deliberate choice as one paradigm in the theory of rational or autonomous action – it is certainly an important phenomenon for any such theory to explain. My argument instead is that it cannot be the fundamental phenomenon, for it is built up from, and at every step involves, the operation of countless non-deliberative processes that are – and must be – quite unlike choice. These processes are not self-aware or reflective, yet they are intelligent and responsive to reasons *qua* reasons. They make us the agents we are, and give our agency its capacity for rational, autonomous self-expression.

To see such processes in operation, we should look first to a different sort of exercise of autonomous agency from paradigm-case reflective choice. We should look to *fluent* agency. To focus on deliberate choice as the core model for what happens when mature individuals act is like focusing on my initial unskilled, very deliberate driving as the core model of what happens when experienced drivers drive, or on the unskilled, very deliberate way beginners attempt to speak or think in a foreign language as the core model of what happens when fluent speakers think or talk in their mother tongue. It is not as if experienced drivers or native speakers do the very same things as the neophyte, only somehow faster and more smoothly. Rather, they do something quite different, and yet equally or more responsive to reasons *qua* reasons, and equally or more self-expressive and autonomous – equally or more *their* actions done for *their* reasons.

For this to be possible, there must exist non-deliberative causal psychic processes “of the right kind” to be aptly responsive to a given consideration as such, and aptly expressive of one’s identity or values, even in the face of competing interests. What would this look like? Consider first an experienced driver, Christine, in a desperate hurry to reach town before the bank closes for the weekend, driving a long gravel country road. She is used to this road, and can drive it at speeds up to 50 mph, despite its many curves and potholes. But driving it that fast, as she is now, requires her full attention, not least because of the risk of equally fast cars moving in the other direction, visible only a split second in advance around the many blind curves. She is on the edge of her seat, muscles taut, intensely studying the road surface ahead to spot potholes and tweaking the steering wheel to weave between them. Rounding a curve onto a straight stretch she suddenly spots an on-coming car, moving slowly at a small distance. In an instant she recognizes that the driver is aged, barely able to see over the steering wheel, and plowing along almost in the middle of the road. The road here is wide enough, however, to allow her to blast past with room to spare. But after a fraction-of-a-second’s hesitation, she instead pulls her foot off the accelerator and brakes, slowing almost to a standstill. She rolls down her dust-covered window and waves broadly to the other driver. He looks up at her. His face, tense from the strain of driving with poor eyesight, relaxes into a smile of acknowledgment, and Christine drives on. Why did she do this, losing precious time? She did not know the other driver, or expect ever to see him again.

Well, the instant she saw the anxious look on the other driver's face, part of her mind empathetically simulated his anxiety, and how it would rattle him were she to surprise him with her speed, blasting past in a cloud of dust. The negative feeling generated by this rapid, non-conscious simulation caused her a moment's hesitation, and drew her attention from her own driving to the other driver's situation. Once that had occurred, her sympathetic feelings responded, "Slow down, back off, make it easy for him." Why were these particular feelings so readily available, and why did they take the form they did? Over the years, Christine had spent time with many older individuals, and knew from experience how, for them, driving could be anxious, tentative, and frightening. She knew that facial expression and driving posture. So as soon as awareness of the other driver's situation had come to the surface, she paused and let up, checking her initial impulse to blast past. Moreover, Christine is someone given to placing people and their feelings ahead of many other concerns or goals. As soon as the thought of relenting occurred, it immediately fitted her feelings.

It strikes me that she acted fully intentionally – she did not act reflexively or simply "find herself" slowing down and wondering why. Instead, she deftly executed a suite of well-timed, well-aimed, and well-coordinated actions. She knew what she was doing, and could accurately answer as to why she did it. Of course, many components of this suite of behavior were not themselves conscious or intended. Her simulation of the other driver's feelings occurred without her will or awareness; the idea of slowing down came not from reflection, but from emotionally colored, lightning-quick association. Only two alternatives entered her mind – speed around him or slow down. Neither this range of possibilities, nor the reasons associated with them, were chosen by her. Her thoughts occurred in quick succession, and she rejected blasting past as soon as the alternative of "backing off" became available. What favored the second alternative was not a "chosen reason," but its fit with her sense of the other driver's fear and vulnerability. Yet the rejection of the first alternative, and selection of the second, surely was her intelligence and agency at work. And the reason that led her to slow down was surely *her* – non-accidentally reflecting her understanding, sensitivity, and priorities.

Her behavior, as described, is perhaps not much different from Aristotle's notion of how the brave soldier might conduct himself in battle or the brave mariner in a storm at sea – acting with speed and deftness, intelligently rather than mechanically, focusing on the essentials,

alive to the state of the world and of the people around him, not locked into a rigid plan, but quickly identifying the most relevant alternative courses of action, seizing upon an appropriate one with a minimum of deliberation, and acting resolutely even when sacrifice is required. According to Aristotle, such an individual is acting in the right way, at the right time, with the right motive, and toward the right end – a model not only of what we have called agent-competence and practical intelligence, but of rational self-mastery and self-direction. We might say that he is an agent *attuned* to the reasons he faces, and *practically so*, since he translates them effectively into action.

I see Christine's less dramatic actions as practically attuned in the same sense, and therefore have no difficulty seeing her action as both rational and autonomous. She mastered both herself and situation, doing so in line with the fundamental springs of her own scheme of values. Despite the lack of reflection or conscious willing, the shaping force for her action came from within, but without self-centeredness. Anyone who knows Christine would agree that, in slowing down and making herself a bit later, she was absolutely being herself.

Why call such action autonomous? After all, it is quite unlike the familiar paradigms of autonomy, in which an agent stands back from her concerns, reflects, and then makes a choice on the basis of a principle she imposes upon herself. My reason is simple. These paradigms themselves are composed of a suite of thoughts and actions, the basic elements of which each must be reasons-responsive but not, on pain of regress, via an exercise of deliberative agency. As we noted earlier, contemporary psychology increasingly points to the role of affect and affective processing in shaping and guiding such coordinated suites of thought and action. Our affective system often responds to perceptual inputs more rapidly than higher cognitive functions, and this initial affective response – positive or negative, as the case may be – then influences subsequent cognitive processing. Affect also mediates such mental processes as attention, association, memory, recall, and execution. In Christine, the negative affect generated by her unconscious simulation of the other driver's situation tended to inhibit her current course of action (blasting ahead), refocus attention, and prompt thoughts of less aggressive alternatives. In her case, time was of the essence. But even the most deliberate of decision-making involves affect in essentially similar ways: by furnishing a "sense" of what needs to be considered, of what evidence is credible, of how weighty a given consideration is, of whether I'm satisfied with the alternatives considered, of when I'm ready to

decide, of how confident I am of my decision, of how motivated I am in enacting it, of how resolute I am in sticking with it, and so on.³⁷

For any decision involving premeditation, then, there must be many unpremeditated, affect- or trust-guided processes underwriting the deliberation and choice, shaping its course, moving it along, giving it force. The self-conscious decider is a captive audience for the considerations that arise in her mind. What these considerations are, as well as how they strike her, will ordinarily not be something she is able to decide. An agent can no more escape such “captivity” to the outputs of her sub-personal self than she can escape epistemic “captivity” to her own sensory and mental experience. Of course, she may be free to ignore any such considerations – but then, where will *that* thought or dismissive impulse have come from? If a deliberated decision is to be autonomous, these non-deliberative elements must be compatible with, even contributory to, autonomy. When they work together in the right way, the result is autonomy – even in the sped-up form found in Christine’s case. Of course, each element need not itself be autonomous – the whole is certainly more than the sum of its parts, and human psychology is rife with such emergent phenomena. But I would insist that the true whole in the case of autonomous action is not the suite of self-conscious mental acts upon which we focus when considering paradigm cases, but rather the whole psyche within which these self-conscious acts are embedded, from which they emerge, and to which they are indebted at every step. *Neurath’s deliberator*, not *Zeno’s*.

Self-constitution through action makes sense only if there is enough of a self already constituted, with sufficiently definite opinions, aims, and ideals, to guide the further self-constitutional process in a meaningful way. It is really self-*re*constitution. In this reconstructive effort, however, the already-constituted self, unchosen as it might be, is *ipso facto* accorded some normative authority to make choices and act. If this cannot be autonomous action, then I cannot be an autonomous actor. Creating the self to which one aspires must be done with the self one has – the self that has this aspiration as well as many elements that may impede it. For example, I do not identify with the tacit racial prejudices that lie within

³⁷ Kant himself, aware of the danger of regress, saw the moral law as coming into the reasoning of the agent through an affect, respect:

Respect (*reverentia*) is, again, something subjective, a feeling of a special kind, not a judgment about an object that it would be a duty to bring about or promote. For, such a duty, regarded as a duty, could be represented to us only through the *respect* we have for it. A duty to have respect would thus amount to being put under obligations to duties (*The Metaphysics of Morals*: 6: 402–403).

me, as they do in most everyone in our country. I wish instead to construct a self unprejudiced in thought and action. But I must recognize that the self with which I am working *is* prejudiced, and I cannot accomplish a rebirth, cleansed of prejudice by endorsing an unprejudiced self-concept. Self-reconstitution is difficult and often only partially successful work. That much self-understanding, at least, and that much willingness to accept that one's acts are the acts of one's "whole" self, like it or not, is essential to winning one's way to autonomy. Deliberation and decision are a small part of our whole intelligent, sensitive, goal-directed psyche. Such rationality and autonomy as we have resides in this whole, not in our self-awareness alone.

A LIST

Musical improvisation, whole-hearted engagement in a conversation or sport, and exercise of a skilled craft to solve a challenging problem are, I would argue, our paradigms of self-directed activity. We see the individual as agent especially clearly in such cases, for it is here that she has the greatest degree of control, self-expression, and attunement to available reasons. Moreover, we are told, individuals characteristically report such activities as *optimal experience*.³⁸ Aristotle would explain: These activities show man acting in accord with his nature, making fullest use of his distinctive capacities. Having such experiences requires, I have argued, not just practical reason and will power, but practical competence, even practical fluency. What might a list of practical competencies look like? Here's a start.

Many general-purpose competencies figure among the conditions for practical competence—e.g. perceptual abilities, inferential capacities, multi-tasking capability, reliable memory, vivid imagination, etc. Others are especially relevant to action and emotional response. We've already mentioned (among others):

- the ability to act intentionally without intending to do so
- the ability to detect, without reflecting and deliberating, conditions calling for reflection and deliberation
- a capacity to direct and focus one's attention in response to cues from one's circumstances, including a capacity to modulate one's attention span and redirect attention, without requiring attention to do so
- the ability to work around or with temptations one cannot suppress

³⁸ See Csikszentmihalyi, *Flow*.

- the ability to reward oneself effectively, accepting life’s diverse pleasures
 - including, say, the pleasure of doing one’s duty
- the capacity of one’s emotions to attune one to demands of a situation
 - e.g. to feel an appropriate degree of fear in response to risk
- the ability to “channel” emotions into effective behavior
- the ability to stick with an intention, other things equal, without constantly revisiting it, along with an opposing disposition to lose confidence in intentions that are not working out
- an ability to represent accurately the mental and bodily states of others or of one’s hypothetical future self
- a degree of self-confidence, including a defeasible self-trust of one’s capacities
- a capacity to be resilient in response to set-backs and not too hard on oneself, with an opposing disposition to accept criticism and the need for change

Here are some others:

- the ability to transfer the motivational interest of an end onto a means, along with a countervailing tendency to lose confidence in ends that come to seem infeasible or to require intuitively unacceptable means
- a degree of self-awareness, and a capacity to self-locate in space and time, sensing oneself as a distinct and continuing being located at a particular causal nexus
- a modulated sense of one’s physical and mental state and capacities – e.g. one’s alertness, capacity for thought and effort, health, arousal, etc.
 - that directly enters into the influences on how one thinks and acts without need for conscious attention
- a felt pressure toward consistency in belief and coherence in desire and comportment, along with a countervailing tolerance for some degree of internal conflict and trial-and-error experimentation

And many others. The list is diverse, but functionally integrated by the requirements for fluent intentional action (much as “administrative skills” or “athletic skills” are functionally integrated by the requirements of certain roles or activities).

Chronic failings in practical competencies can be very serious for an agent – often more serious than garden-variety irrationality. Certain familiar psychological disorders can be seen as persistent deficits or hypertrophies of practical competencies: Chronic inability to maintain focus on a task (attention deficit disorder), chronic and inflexible inability

to redirect attention from a fixed concern (obsession), chronic inability to attach affect to goals or gain reward from accomplishment (depression), chronic inability to regulate motivation by affect or thought (addiction, compulsion), chronic inability to resist temptation (e.g. impulsivity disorder, kleptomania, etc.), chronic inability to sense the emotional states of others non-inferentially (sociopathy, autism), chronic arousal of the self-schema (narcissism), and so on. It may not be appropriate to group the behavior of individuals suffering from such disorders with the behaviors typical of garden-variety irrationality – losing one's temper, failing to match means to ends, succumbing to peer pressure or a sales pitch against one's better judgment, giving in to the temptation to overeat, etc. What is needed in their circumstances is something other than sobering up, calming down, thinking things through more fully before acting, weighing options more carefully, or exerting more "will power."

SOME BENEFITS

I do not need to shill for the importance of cultivating the capacities I have grouped under "practical competence" – self-help books abound with this sort of thing. Similarly, tracts in analytic psychiatry and cognitive therapy can do much better than I at categorizing and anatomizing failures of these competences and showing the toll they can take on one's personal efficacy and well-being. Research psychologists are amply critical of the partiality and inaccuracy of our first-personal self-narratives and self-knowledge claims, and the mischief that results. These many sources point to the fundamental role of the "larger self" – underneath and around the self-aware deliberative self – in shaping what we think and do, and to the self-deluding character of the thought that we can transcend or master it with a "rational will." Yet although we cannot master this larger self – Who, after all, is supposed to do the mastering? – we can better understand it. And we can use this understanding to help discern and develop the forms of mastery of which we *are* capable.

I suggested at the outset that there might be some benefits in the theory of practical reason to thinking about agency along the lines laid out in this chapter. What might they be?

One benefit could be a rethinking of the "internalism"/"externalism" debate about the content of moral judgments. In the existing debate, all sides accept that there appears to be a non-accidental tie between making a moral judgment, such as "I ought to *X*," and possessing some

motivationally-ert pro-attitude toward *X*-ing. Judgment internalists take their cue directly from this datum, and give a non-factualist reading of such judgments according to which expressing such an attitude is their primary function. For them, motivation is internal to the content of moral judgment. Externalists, who insist that moral judgments are factual, functioning primarily to state normative propositions with ordinary truth-conditions, cannot build motivation into their content. Instead, they have been forced to give more indirect explanations of the intimate link between judgment and affect. Often these are criticized as not intimate enough.

The dual-process models characteristic of contemporary cognitive psychology promise a way of explaining the psychological reality of internalist intuitions without abandoning factualist semantics. We have inherited from countless generations before us a remarkably effective sub-personal learning and evaluation system, which triggers a rapid affective response directly to perceptual experience or mental suggestion. This response then primes and cues judgment while simultaneously inducing emotional response and corresponding motivation or act-preparedness. Does this point to judgment internalism, even emotivism?

Not at all. It's worth keeping in mind that the early elicitation of affect is characteristic not only of moral judgments, but also our responses to images of spiders and smiling babies, even to lists of ordinary descriptive terms ("smooth" elicits positive affect, and therefore finds its way into countless advertisements; "jagged" is not, and therefore does not). "Affective coding" is therefore *normal* for judgments, including paradigmatically factual judgments ("The surface is smooth"; "The surface is jagged"). Small wonder that judgments containing moral terms, which have powerful positive and negative affective valence ("good," "kind," and "right" vs "rotten," "cruel," and "wrong"), have a non-accidental association with positive or negative affect and motivation.

This fact helps explain why judgment internalism *seems* so plausible. Subjects shown a drawing depicting a man menacing a child show a negative affective response milliseconds before the sensory signal has reached the higher cortices, where such verbal labels as "wrong," "bad," or "assault" are available and triggered. Thus an "anti-attitude" seems to be *part of the very thoughts* WRONG OF BAD OF ASSAULT. But seeing a spider image, too, elicits a negative affective response in the first milliseconds, even before the thought SPIDER is cued. Should we say that an anti-attitude is part of the concept SPIDER as well? Or that a pro-attitude is part of

the concept BABY? Should we say that these, too, are normative concepts or “action-guiding” concepts? A more promising line would be that because the judgment and the felt pro-attitude have a common cause, stemming from the pre-conscious affect-coding, there will be an appearance that the emotive attitude is “part of the content” of the judgment. The two will be seamlessly co-present in the conscious mind. But once we see that this does nothing to distinguish moral judgments from many prosaically factual judgments, the pressure to adopt internalism with respect to the content of moral judgments is removed.³⁹

This conclusion can be strengthened by noting that the dual-process explanation accounts for another intuitive datum. Some moral judgments are situational, stimulated by a concrete, salient description or perception. Here the immediacy of affect is palpable. But there are also highly theoretical moral judgments. For example, an advisor to a legislator might be asked to determine what rate of progressivity in income tax is morally appropriate, taking into account trade-offs with economic growth, wage disparities, and effects on human capital accumulation. Or one might attempt to determine just what portion of an individual’s income or effort should be devoted to charitable causes, taking into account differences in income and wealth, reasonable social expectations, etc. When one reaches a conclusion in such a judgment there is little or no affective immediacy. They do not appear to have motive force “built into them.” More likely, having reached such a judgment, one might have to cultivate some affective confidence in it or internal motivational support for it. For example, we think a utilitarian insincere who, after performing an elaborate cost-benefit analysis, announced a conclusion about optimal tax schedules or charitable giving without himself experiencing any accompanying motivational force. Haidt’s work on a dual-process model of moral judgment revealed that the priority of affect in moral judgments of scenarios is largely absent when the scenarios are presented in such a way as to make abstract moral reasoning more salient than a one-off moral judgment.⁴⁰ An externalist can give a straightforward explanation of these data: there is no built-in motivational content in moral judgments, rather, the affective pathway to moral judgment is differentially triggered by concrete vs abstract moral judgments. At first blush, judgment internalism, which imputes built-in motivational content in all moral judgments,

³⁹ At least, this source of pressure is removed. There are others.

⁴⁰ Murphy *et al.*, “Moral Dumbfounding.”

fits these data less well. So, one benefit of thinking about the puzzles of practical reason in terms of a dual-process model of human psychology could be rethinking whether the intuitive data really support judgment internalism after all.

Moreover, we now are in a position to see a possible second benefit. Recall the contestation, mentioned at the outset, between internal and external perspectives in the theory of action and practical reason – between a first-personal approach “from the inside out” and a third-personal approach “from the outside in.” Of course, both sorts of understanding will figure in any complete account. But there is something more to say: the usual way we talk about internal and external is misleading. Given the limited, and sometimes marginal, position of the self-aware, deliberative “self” within the larger selves by means of which we are competent agents and intelligent responders to reasons, we cannot equal a first-personal approach with the standpoint of the portion of the self that happens to be conspicuous in immediate awareness or open to introspection. “I” refers to the larger self as well, and much that is distinctively me and equips me for practical rationality resides in the whole. Thus an approach working “from the inside out” must recognize that the self-conscious agent needs to supplement introspection with an explanatory, third-personal perspective to understand just what he is doing, and why. That is, the true perspective of the self is available only first- *and* third-personally.

A third possible benefit concerns the much-vexed question of the normative status of means-end coherence. We tend to be of two minds about this. On the one hand, it seems to be a mark of practical rationality to identify means appropriate to one’s ends, and to be able to transfer motivational force from the ends to the means. On the other hand, when someone has manifestly irrational or horrific ends, or ends that require horrific means, we do not think it would be a mark of rationality for him to seek out and pursue the requisite means. The only rational response to certain ends or means might be to abandon the end altogether. As Hare once wrote,⁴¹ one can be lulled by means-end statements like:

- (1) If you want sugar in your coffee, you should use a spoon.
- (2) If you want your coffee warm, you should heat the jug.

⁴¹ See R.M. Hare, “Wanting: Some Pitfalls,” in R. Binkley, R. Bronaugh, and A. Marras, eds., *Agent, Action, and Reason* (Toronto: University of Toronto Press, 1971): 81–97.

into thinking that these are “hypothetical imperatives” of rationality—that is, that there is a rational principle of the form “If you want X , and Y is the most effective means to $X \rightarrow$ you should Y .” But then notice:

- (3) If you want to commit murder on a massive scale, you should poison the city water supply.

Suppose that (3) were accepted as a hypothetical imperative of the form “If you want to commit murder on a mass scale, and poisoning the city water supply is the most effective way to do this \rightarrow you should poison the city water supply.” Then if someone did want to commit murder on a massive scale, he could perform *modus ponens*, detach the consequent, and claim that rationality commends his poisoning the city water supply. The following is truer to our rational convictions:

- (4) If you want to commit murder on a massive scale, you should see a psychiatrist at once.

Note that (4) does not counsel means-end coherence – going to a psychiatrist is certainly not the most efficient way to commit mass murder. Nor need one withdraw (4) if this individual has no other aim that a trip to the psychiatrist would serve. Rationality, we might think, can tell against an entire system of ends, however coherent.

Still, there is an evaluative difference, irrespective of the nature of the end or the means, between an agent who shows ingenuity and success in matching means to ends vs an agent who elects ill-considered and ineffectual means, or who fails out of self-indulgence to follow through. What’s wrong with the second agent is that he is manifesting *agent-incompetence*, which is a distinct defect even in those cases, like wanting to commit mass murder, in which defective agency would on the whole be morally fortunate. Just as we don’t want to confuse “technical imperatives” like (1)–(3) with rational imperatives, we don’t want to confuse *rational criticism* of an agent’s ends or means with *technical criticism* of his agent-competence. Even as we are horrified by the countless lives his actions cost, and the wreckage he finally made of the ideals of the Revolution, we cannot but admire the fluent agency and practical intelligence of Napoleon in battle. Mere possession of an aim – a desire, a whim, an obsession, an ambition – may not suffice to make it rational to pursue. But *any* aim suffices to provide an opportunity to manifest fluent agency and practical intelligence. Or fail to.

The most common way by which we competently match means to ends is not via practical reasoning, but via learned *habitus*, trial-and-error,

imitation, or “intuition.” Entering a room, I flip on the light to see what I’m doing, place one foot in front of the other to keep moving, turn my head toward the table, focus my gaze on the unopened letter on the table, extend my arm while opening my fingers in order to grasp and examine it, etc. In this little scene, means are being matched smoothly to ends, but I rather doubt we should see this sort of learned coordination of thought and act as the work of practical rationality. After all, animals perform essentially the same sorts of purposive, non-deliberative, semi-habitual means-ends matching. Your cat, hearing the door open, runs over to rub against your leg and nudge you gently toward the kitchen where the food is kept. A seeing-eye dog leads her master smoothly through a crowded room, successfully tracking the trajectory of dozens of people and deftly guiding her master through the constantly changing gaps between clumps of people. If it seems inappropriate to credit these actions on the part of cats and dogs to practical rationality, then perhaps we should say the same about our own like actions.

Humans, of course, differ from many other animals in their ability to engage in what is identifiably practical reasoning in matching means to ends—for example, when they conceptualize alternative possibilities, assess and compare costs and benefits, and form complex, branching contingency plans. This capacity equips us for the spectacular explosion in the range of possible ends and means that the human race has brought to the world, for better or worse. Perhaps it will even save us from ourselves, and save life on this beleaguered planet from self-destruction. For practical reasoning is the only capacity on earth able to translate a vivid representation of remote future possibilities into a apt plan for action today, and to guide its implementation across the decades to come.

But it does not take practical rationality to be good at matching means to ends. This had better be so – calling upon the deliberative capacities of practical reason is *itself* applying a means to an end, and if *this* required practical deliberation, then . . . , and so on. With lots of experience and a little luck, we can learn to deploy this distinctive means on those occasions that warrant it, and to carry it off reasonably well without overdoing it. That is, we may become practically competent – perhaps even fluent – at practical reason.