From A Cognitive Theory of Consciousness. (B.J. Baars, 1988, NY: Cambriddge University Press).

Chapter Seven

Model 5:

Volition as ideomotor control of thought and action.

"We know what it is to get out of bed on a freezing morning in a room without a fire, and how the very vital principle within us protests against the ordeal. Probably most persons have lain on certain mornings for an hour at a time unable to brace themselves to the resolve. We think how late we shall be, how the duties of the day will suffer; we say,"I îmustï get up, this is ignominious," etc; but still the warm couch feels too delicious, the cold outside too cruel, and resolution faints away and postpones itself again and again just as it seemed on the verge of bursting the resistance and passing over into the decisive act. ...

Now how do we îeverï get up under such circumstances? If I may generalize from my own experience, we more often than not get up without any struggle at all. We suddenly find that we îhaveï got up. A fortunate lapse of consciousness occurs; we forget both the warmth and the cold; we fall into some revery connected with the day's life, in the course of which the ideas flashes across us, "Hollo, I must lie here no longer" --- an idea which at that lucky instant awakens no contradictory or paralyzing suggestions, and consequently produces immediately its appropriate motor effects. ...

It was our acute consciousness of both the warmth and the cold during the period of struggle, which paralyzed our activity then and kept our idea of rising in the condition îwishï and not îwillï. The moment these inhibitory ideas ceased, the original idea exerted its effects.

This case seems to me to contain in miniature form the data for an entire psychology of volition. ..."

--- James (1890, Vol. II, p. 524-5).

‹b ‹

7.0 Introduction.

We begin our chapter on volition with the image of William James on a cold winter morning, reluctantly trying to persuade himself to get out of bed. For James, this image goes to the very heart of the psychology of volition. He believed that a successful act of will does not typically emerge from some titanic inner struggle. Rather, he claims, we simply wait until the conscious image of the action can emerge for some time without competing images or intentions. At that moment the action occurs automatically, spontaneously, and without struggle.

We will first consider whether there is a problem of volition at all. To answer this question we seek contrasting pairs of actions that differ only in that one action is voluntary and the other, involuntary. These empirical contrasts can constrain theory, just like the contrasts between conscious and unconscious events (1.x). This evidence indicates that the issue of volition is very real indeed. That is to say, the voluntary/involuntary contrasts highlight core psychological issues, such as automaticity due to practice, errors in speech and action, and psychopathology. Further, we can borrow James' solution to the problem of volitional control, and interpret it easily in terms of global-workspace theory.

James explains conscious control of action by an îideo-motor theoryï in which conscious goal images without effective competition serve to organize and trigger automatically controlled actions, which then run off without further conscious involvement. For James, conscious contents are inherently "impulsive"; everything else is automatic. The only conscious components of action are:

(a) the "idea" or goal-image (really just an image of the outcome of the action);

(b) perhaps some competing goal-image;

(c) the "fiat" (the "go signal");

and finally,

(d) sensory feedback from the action.

In this chapter we see how GW theory invites a natural interpretation of James' ideomotor theory. One use of the GW architecture is to have multiple unconscious systems inspect a single conscious goal, and to compete against it if it is inadequate. That is to say, the architecture allows multiple‹j ‹ unconcious criterion systems to îmonitorï and îeditï any conscious goal or plan. This implies that any conscious goal image that is conscious long enough to succeed in recruiting and executing an action has been tacitly edited by multiple criteria, and indeed we claim in this chapter that voluntary action îisï tacitly edited action. Conversely, involuntary actions, like slips of the tongue, are actions that îwould have beenï edited and changed, îifï there had been enough time and capacity for unconscious editing systems to be brought to bear upon the conscious action plans. This conclusion has a wealth of implications for the understanding of unintentional acts found in slips, automatisms, and psychopathology. It even suggests a theory of hypnosis and other "absorbed" states of mind, in which there is minimal editing of conscious events.

Of course William James himself could not speak in these terms, because of his resistance to the notion of unconscious mental processes (1.x). But his ideas make perfectly good sense in modern cognitive garb.

Before we proceed to develop these ideas, it is useful to be clear about the issue of volition itself. Many behaviorists and others have claimed that there is no problem of volition at all. What evidence do we have to the contrary, that "the will" matters?

7.1 Is there a problem of volition? Some contrasts between similar voluntary and involuntary actions.

With the rise of physicalistic psychology at beginning of this century, many psychologists tried to make the case that there really is no question of volition, just as there was no true scientific issue of consciousness (e.g. Watson, 1925; Razran, 1965). Behaviorists and other physicalists at first believed that any apparently "voluntary" action can be reduced to a chain of conditioned, simple, îphysicalï reflexes. Later, when reflexes proved too simple and rigid, the unit of behavior was generalized to other stimulus-response relationships, but the goal still remained to reduce voluntary, goal-directed actions to simple, physical input-output relations (viz., Baars, 1986a). This was thought to eliminate any scientific question of volition once and for all. Was there any truth to this claim? Is there indeed a scientific question of volition?

For an answer we can look to pairs of actions that appear similar on the surface, but which differ in respect to volition. That is, we can carry out a contrastive analysis on the issue of volition, just as throughout this book we have contrasted comparable conscious and unconscious events (1.xx). This is helpful not just to answer the questions about volition raised by Pavlov and Watson --- it also defines major constraints to be‹j ‹ satisfied by any theory of normal voluntary control. Any such theory should be able to explain why, of two similar-seeming actions, one seems to be voluntary and the other not.

There are some obvious examples of such contrastive pairs, as we see in Table 7.1.

Table 7.1

Closely comparable contrasts involving volition.

Involuntary	Voluntary

îNon©voluntary ïîaïîutomatisms: wanted, but not controllable in detail.ï

Automatic components of normal actions.	The same actions before automaticity.
Reflexes.	Purposeful imitations of reflexes.
Actions controlled by brain stimulation of the motor cortex.	The same actions initiated by the patient.
Autonomic functions. (heart rate, peristalsis, skin conductivity, etc.)	Autonomic functions under temporary biofeedback control.
Spontaneous emotional facial expressions (Ekman, 1984)	"Social" expressions.

îCounter©voluntary automatisms: unwanted. ï

Slips of speech and action.	Purposeful imitations of slips.
Pathological symptoms: out©of©control actions, images, inner speech, and feelings.	Purposeful imitations of symptoms.
	Regained voluntary control after "practicing the symptom."
Voluntarily resisted habits (e.g. unwanted habits)	Voluntarily controlled automatisms.

îNon©voluntary vs. counter©voluntary events.

ïNotice first of all, that the "involuntary" events listed on the left side of Table 7.1 are of two kinds. First, îautomatic ïprocesses are part of every voluntary act, and while people cannot control them in detail, they are perceived to be consistent with our goals. We want them. A skilled typist does not control each finger movement in detail; a skilled reader does not perform letter identification consciously, etc. Yet because automatisms serve our voluntary goals Table 7.1 calls these "non™voluntary automatisms." On the other hand there are clearly îcounter©voluntary actionsï such as slips of the tongue. Here, too, there are automatisms at work, but they are perceived to be out of control, unwanted, against one's will. These two kinds of involuntary action may be closely related; for example, any automatic component can become counter©voluntary simply by resisting it. We can look at a word on this page, thereby triggering automatic reading processes that are not under detailed voluntary control. This wanted automatism can become counter©voluntary simply by resisting it. Thus we can try to resist the act of reading after looking at a word, or a knee©jerk reflex after striking the patellar tendon. In this way any automatism can be made counter©voluntary. This close relationship between non©voluntary and counter©voluntary actions makes it useful to consider both under the rubric of "involuntary" activities. Whenever there is a possibility of misunderstanding, we will choose an unambiguous term like "automatic" vs. "counter©voluntary."

That being said, we can go on to discuss Table 7.1.

îSlips of speech and actionï

Imagine repeating a slip of the tongue you have just made. The slip itself is experienced as involuntary; its imitation is voluntary. And yet the two isolated actions are much the same as far as an outside observer is concerned. Some famous slips by A.W. Spooner illustrate the point:

(1) Instead of "our dear old Queen" --- "our queer old Dean."

(2) Instead of the hymn, "Conquering Kings their titles take." --- "Kinquering Congs their titles take."

(3) Upon dismissing a student, he intended to say "You have deliberately wasted two terms, and you will leave by the down train" --- but actually said, "You have â ‹j
‹
deliberately tasted two worms, and you will leave by â the town drain."

Let us suppose Reverend Spooner actually made these slips (there is some doubt: see Fromkin, 1980). Now imagine that Spooner îrepeated each slipï immediately after making it, as exactly as possible, so that it was said again by the same speaker, in the same tone of voice, at the same speaking rate, and so on. What is the difference between the slip and its voluntary repetition? Surely there is no basic îphysicalï difference, nor any real linguistic difference. The main difference is psychological. In the first case, the utterance was involuntary and unwanted; in the second, it was voluntary (Baars, 1985; in press).

But what a difference this invisible difference makes! In the first case, the speaker fails to execute his intention. If he becomes conscious of his error, he will experience îsurpriseï at his own utterance. Now we can observe the whole panoply of physiological reactions that make up the Orienting Response (1.x). He may be embarassed and apologetic. Having failed to carry out his intention, he may try again. If, like Spooner, he is also head of one of the Cambridge colleges he may become a figure of fun in student folklore. If he makes involuntary errors so often that he can no longer function effectively, he may lose his position, be examined for neurological problems, etc. None of these consequences follow from doing physically identical imitations of these slips, if they are voluntary. If Spooner were voluntarily making the slip to amuse his audience, or if someone quotes a slip in a discussion of voluntary control, none of these consequences follow; nor is the speaker likely to be surprised by the "slip."

Thus two identical actions may be psychologically quite distinct, but not because of a difference in complexity, as the early behaviorists thought. Voluntary actions are not just complicated agglomerations of simple reflexes. Involuntary components put together do not result in a voluntary act. Something else is involved in volitional control. Consider two more contrasts of this kind.

îThe loss of voluntary control with practice.ï

It is easy to see a voluntary act transformed into an involuntary one: we only need to practice it to the point where most of it fades from consciousness (5.xx). We have previously pointed to experiments in which predictable skills are highly overlearned, and which generally show a loss of voluntary control (La Berge, 1984; Shiffrin & Schneider, 1977; Sternberg, 1963).‹j ‹å

All actions have involuntary components. Most details of routine actions like reading or writing must be automatic: we could never control their numerous details, given the limited capacity of the conscious and voluntary system. Usually only the novel features of an action are conscious and under voluntary control (x.xx) (Reason, 1984). But non©voluntary automatisms can sometimes become unwanted or counter©voluntary.

This becomes clear when we try to control "bad habits" that have been practiced for years: almost everyone seems to have at least one, whether it is over-eating, smoking, nervous gestures, etc. These habits are characteristically difficult to control voluntarily; they escape control especially when conscious attention is directed elsewhere. No doubt unwanted habits have multiple causes, but it is easy to demonstrate that sheer automaticity makes it hard to stop an action once its normal triggering conditions are given. As we pointed out above, looking at a word îwithoutï reading it seems to be quite impossible (viz., Shiffrin & Schneider, 1977; La Berge, 1984). The very act of looking at printed words seems to trigger automatisms; To block them we must look away, degrade the word visually, or perhaps focus on only a fraction of one letter. Sternberg's well-known experiment in automatic memory search makes the same point (Sternberg, 1963). The subject really cannot stop the search process when the target letter is found; it just runs on to the end of the memory set (see secton 1.xx). Reason (1983, 1984) has presented a detailed analyses of catastrophic accidents showing that many of them may be due to hard-to- control, highly practiced automatisms that were triggered out of context at the moment of the accident. Several of these accidents led to the death of the person making the error --- about as strong an argument for the involuntary nature of automatisms as we might wish to have.

îInvoluntary automaticity involves a loss of conscious access.ï

Loss of voluntary control over details of an action seems to follow a loss of conscious access to the details. Langer and her co©¨workers have conducted some elegant experiments to support this point (e.g., Langer & Imber, 1979). These investigators were pursuing the hypothesis that perceived competence affects one's performance: the more skilled we think we are, the better we perform --- providing that we cannot monitor our performance directly. One way in which we lose touch with our own competence

is by automatization; when we become skilled readers, musicians, or truck drivers, we lose conscious access to many details of our own actions, and hence become more vulnerable to false‹j ‹ attributions about our own performance. This line of reasoning led Langer and Imber (1979) to devise a simple coding task that people could learn to the point of automaticity in a matter of minutes. Letters of the alphabet were to be recoded into a two-symbol code;
the letters A©¨I were "a triangle plus the înïth letter after A"; letters J¨©R would be "circle plus the înïth letter after J," etc. Thus the letter "B" would be "triangle plus 2", "L" would be "circle plus 3," etc. A preliminary group of subjects reported that they were still conscious of task details after recoding two sentences; after six sentences, they were no longer conscious of the steps. The task had become automatic.

Langer and Imber (1979) now compared the effects of conscious access and automaticity. A Moderate Practice group recoded only two sentences, reported being conscious of details, and was able to specify more steps in the task than the High Practice group, which recoded six sentences and reported automaticity. Now Langer and Imber devised an arbitrary task in which some of the subjects would be called Bosses, others were called Assistants, and a third group received no label. In fact, the three groups did the identical task; the assumption was that the labels would affect the self-confidence of the subjects. Afterwards they were asked to do the coding task once again. "Bosses" performed much as before, no different from the No Label group. But "Assistants" now performed much worse ¨¨¨ îifî the coding task was automatic. "Assistants" who were highly automatic in the coding task made four times as many errors as before, and took 40% longer to finish. In the Moderate Practice condition, where the coding task was not automatic and consciously accessible, "Assistants" did as well as "Bosses."

The simplicity and effectiveness of this study is quite remarkable. And the interpretation is quite clear: if we have no conscious access to our own performance, and if some reliable source of information îseemsï to indicate that we are doing quite badly, we tend to accept misleading feedback because we cannot check our own performance. With direct conscious access to our performance we are much less influenced by misleading labels. These results suggest that three things go together: losing voluntary control over action details, losing consciousness of them, and losing the ability to îmonitor and editï the details. Indeed, the ability to monitor and edit a planned act may be the essence of voluntary control (7.32).

While we may speak of "conscious" monitoring and editing, the fact is, of course, that we are generally not conscious of the rules and criteria by which we do our monitoring. If we find a syntax error in inner speech, we do not consciously say, "Aha! lack of number agreement between noun and verb!" Not even linguists do that. Rather, we simply "know" immediately that the conscious plan is in error. The rule systems that spot the error are quite silent in their details. Thus it is not consciousness that îdoesï the monitoring and editing; rather, conscious experience of the event îfacilitatesï editing and monitoring by‹j ‹ unconscious rule systems, just as the GW architecture facilitates the ability of many specialized processors to review a global message.

Thus any complete theory of voluntary control must explain the automaticity dimension: why, with practice, we lose both conscious access to and voluntary control over the details of an action.

îPathological loss of voluntary control.ï

Psychopathology is the study of repeated, dysfunctional errors that are generally îknown to beï errors by the person who makes them --- "slips" of action or experience that escape attempts to control them, over and over again. Almost without exception, psychopathology in the neurotic range involves a loss of voluntary control over inner speech, feelings, mental images, or overt actions. Loss of control over îinner speechï is a factor in obsessive or delusional thinking, and in some auditory hallucinations; out of control îbodily feelingsï play a role in pathological anxiety, conversion hysteria, and depression; uncontrolled îmental imagesï are at the core of phobias; and when îactionsï run out of control we find compulsive or impulse-control pathology.

We can illustrate all these points with a single patient who suffered from a variety of symptoms. Consider Anna O., the classical early patient of Breuer and Freud (18xx), who suffered From a very severe case of conversion hysteria. As Erdelyi (198 describes the case (p. 20),

"Anna O. became Breuer's patient in 1880 at the age of 21 when, under the pressure of nursing her dying father, she suffered a nervous collapse. She developed a veritable museum of symptoms which included a labile (variable) pattern of incapacitating paralyses of the limbs; depression and listlessness; terrifying hallucinations of snakes, which transmogrified into death's heads and skeletons; painful coughing fits, especially in reaction to music; a period of severe hydrophobia, during which she could not bring herself to drink water; amnesias (blackouts) for recent events; a blinding squint; severe paraphasia (loss of language ability); anorexia (unwillingness to take food); and several other serious dysfunctions."

It is the îloss of desired control ïthat makes these symptoms pathological. Not moving one's limbs is quite all right if one doesn't want to move them; depression and sadness due to a loss is quite normal; strong squinting is a good idea in the middle of‹j ‹ a sun-drenched desert; even images of snakes and death's heads can be quite normal for a reader of Gothic fiction (after all, thousands of people voluntarily go to horror movies or read Gothic tales); even amnesias for recent events can be normal when we want to deliberately forget or ignore them. These events become pathological when people do not want them. Those who suffer from these symptoms try hard and often to master the involuntary feelings, thoughts, actions, or images, but they fail over and over again, in spite of often desperate efforts (e.g. Horowitz, 1975 ab). It is not the îcontentï of the thoughts, feelings, and actions that is problematic: it is their occurrence out of an acceptable context, out of the perceived control of the sufferer. Thus the issue of voluntary control is at the very core of human psychopathology, and an understanding of psychopathology must be grounded in an adequate theory of volition (see 7.8).

There is a clinical intervention that is sometimes very effective, which seems to act directly on the mechanism of voluntary control. This paradoxical technique is called "negative practice," or "practicing the symptom" (e.g., Levine & Scheff, 1980). If a person has a specific phobia, he is told to voluntarily bring forth the fearful images and thoughts; if he is a stutterer, he is to try stuttering voluntarily whenever he stutters spontaneously; and so on. Although this technique has been known for decades, it has only recently begun to be systematically tested in a variety of problems. Some of the results are quite remarkable. Children who have stuttered for years are told to stutter deliberately for 30 seconds each time they do so involuntarily. As a result, they sometimes stop stuttering in a day or two, with a 75% success rate (Levine and Scheff, 1980; Levine, Ramirez, and Sandeen-Lee, 1982). There are many cases in which the paradoxical technique works remarkably quickly to stop anxiety attacks, compulsive actions, tics, involuntary phobic images, La Tourette symptoms, and the like. Here is a case where counter©voluntary automatisms are turned into wanted but non©voluntary automatisms, just the opposite of the case of "bad" habits discussed above. Of course, "practicing the symptom" is not a cure-all. But it has been reliably observed to stop pathological symptoms with remarkable speed, often after years of helpless struggle.

Of theoretical interest here is the neat contrast between voluntary and involuntary control in the paradoxical technique. A habitual stutterer has typically struggled thousands of times against the tendency to stutter. This repeated attempt to exert voluntary control rarely works. The paradoxical intervention requires him to stutter deliberately, îto do voluntarily what normally happens involuntarilyï --- and rather magically, in many cases the problem disappears. One fascinating possibility is that the paradoxical intervention (which is, after all, only a switch in the direction of voluntary effort) operates through the voluntary control system. If this is true, then it may be that the symptom itself is an error in voluntary control. Much psychopathology may involve "errors of the will." These‹j ‹ speculations pose some important questions. We will return to them when we attempt to model the voluntary-involuntary contrasts of Table 7.1 (see 7.xx).

îVoluntary action is consistent with one's dominant expectations.

ï

The cases discussed above --- slips, automaticity, and psychopathology --- suggests that îcounter©voluntary action always surprises the actor. ïThis is also true for non©voluntary automatisms like reflexes when we resist them. Thus any non™voluntary automatism is either surprising, or can be made to be surprising when it is resisted. Under these circumstances the automatism îïviolates dominant expectations (the dominant context hierarchy) (viz., Baars, in press) º(Footnote 1). Conversely, îvoluntary action ïseems always to be consistent with one's dominant expectations.

There is direct evidence for this proposal from the study of slips in speech and action (Baars, 1980 and in press; Reason, 1984). First, of course, we know that people often express surprise when they make a slip. The Galvanic SKin Response (GSR) is a well©established measure of surprise. GSRs monitored immediately after experimentally elicited slips are quite large when a sexually explicit slip is made, much smaller when a neutral control slip occurs, and non-existent when the subject makes a correct response (Motley, Camden & Baars, 1982). Thus the more surprising (dominant context©violating) the slip, the larger the GSR. Further, many slips are spontaneously self-corrected immediately after they are made, again suggesting that they surprise the speaker (Baars & Mattson, 1981). This evidence supports the idea that îcounter©voluntary action violates the expectations of the actorï, even when an outsider might not notice anything unusual (Baars, in press, d).

îSome neurophysiological observationsï

Neuroscientists have never stopped using words like "voluntary" and "involuntary" to describe some obvious phenomena. Reflexes are obvious examples of involuntary actions; so are autonomic functions like peristalsis, heart rate, sweat gland activity, and the like, as opposed to the control of skeletal muscles, which is voluntary in the usual sense. We now know that autonomic functions can come under voluntary control at least temporarily when people are given conscious feedback signals activated by the autonomic functions (2.xx, 3.xx). Biofeedback‹j ‹ training seems to bring autonomic responses under the control of the conscious/voluntary system. All these cases present obvious contrasts between voluntary and involuntary control of the same physical functions.

Another remarkable example of a neurophysiological contrast between voluntary and involuntary control is cited by Penfield and Roberts (19xx). These neurosurgeons used a low-voltage electrode to explore the exposed cerebral cortex of conscious patients, in order to identify and avoid critical areas where surgery might cause serious damage. In one case as the surgeon probed the motor cortex, the patient's hand moved, and the patient was asked, "Are you moving your hand?" --- whereupon she replied, with perfect accuracy, "No, doctor, îyouï are moving my hand." How could the patient possibly tell the difference between the brain mechanisms that were under "her own" versus the surgeon's control? We do not know, of course, but her ability to make this distinction suggests that there is a major difference between voluntary and non-voluntary control.

In sum: Is voluntary control really a psychologically significant issue? Facts like these indicate that it is indeed. From here on we will assume that common sense is well justifi in giving volition a fundamental psychological role (Footnote º2).

7.2 Voluntary action resembles spontaneous problem-solving.

In Chapter Six we worked out a way of understanding the Conscious-Unconscious-Conscious (CUC) triad found in so many types of problem solving. Thus, in answering a question we are conscious of the question in detail, but not of searching for the answer, though the answer is again conscious (6.xx). In creative problem solving we are aware of the type of solution we need, but not of the incubation process which eventually brings it to awareness. And so on. Further, we have addressed the whole question of what is meant by an îintentionï by considering the tip-of-the-tongue (TOT) state, concluding that even as we are searching for the right word, there is a state of mind which constrains the search, which constrains limited capacity, but which does not have qualitative conscious contents like color, texture, or flavor. This "intention to say so-and-so" was called a dominant goal context (6.xx).

Voluntary control resembles spontaneous problem-solving in many ways. As James suggests, in voluntary action a conscious goal image may be carried out unconsciously, and the results of the action often become conscious again (7.xx). For illustration, we will ask the reader to turn this book upside-down. (It is helpful to actually carry out this little experiment in‹j ‹ self-observation.) Clearly the reader is conscious of the request to turn the book upside-down, and perhaps of some visual image of how this might be done. However, the request is ambiguous: is the book to be turned in the horizontal or the vertical plane? This ambiguity may be conscious for some readers and unconscious for others. The mechanics of controlling hand and arm muscles are surely not conscious, although îchoice-points and obstaclesï (how do I turn the book upside-down without spilling my coffee?) may be conscious. And of course the results of the action will be conscious.

Further, there is a set of îconstraintsï on the action, represented in GW theory by the îdominant goal contextï, which are not likely to be conscious at any time during the action (Figure 6.x). We probably turn the book over with maximum economy of movement, rather than sweeping through the air with graceful, elaborate gestures. Then there are constraints imposed by the need to maintain physical control of the book; we are not likely merely to flip it up into the air and let it fall helter-skelter. Even further, there are constraints of convenience, such as keeping track of one's place even while indulging in this little thought experiment. We must stop reading while the book is being moved, and we make automatic postural adjustments to balance the changing forces on the body. Finally, there may be social considerations --- if we are in public, is anyone watching our peculiar behavior? While some of these considerations may be momentarily conscious, many of them will be unconscious, but they still serve to constrain the action.

In a real sense the action that results from this complex set of momentary conscious and unconscious constraints is a îsolutionï posed by îproblemsï triggered by the conscious goal, and bounded by numerous physical, kinetic, social, and other contextual considerations. It makes sense therefore to treat voluntary control as a kind of problem-solving (6.0).

7.21 Cooperating automatic systems control most of a normal "voluntary" action.

The bulk of spontaneous problem-solving is unconscious (6.xx). The same is surely true of voluntary actions. Much of our intention to perform a particular act must be formulated unconsciously, and the muscular effectors and subgoals needed to carry out the intention are also largely unconscious. Thus many systems cooperate in creating a voluntary act. It is good to keep this great amount of cooperative processing in mind during the coming discussion, which will focus mostly on the îcompetitiveï aspects of voluntary control.

Notice, by the way, that the same systems may cooperate most‹j ‹ of the time, only to begin competing when the action runs into trouble. If many systems work together to structure normal speech, a slip of the tongue will seem erroneous to some but not all of those systems. When Spooner slipped into "our queer old Dean," he made no error at all lexically, phonetically, syntactically, or even in pronunciation rules. The only systems able to detect the errors are semantic and pragmatic: that is, the systems that control meaning and communicative purpose. Those are the only levels violated by the slip. It would seem to follow that those systems begin to compete against the error, while the others continue to cooperate.

7.22 We become conscious of underdetermined choice©points in the flow of action.

If we are unconscious of these routine, cooperating systems, what are we conscious of? Our previous discussion (5.xx) suggests that the most informative aspects of action should be conscious: that is, those that are unpredictable and significant. It is the îunderdetermined choice©points ïin the flow of action that should be conscious most often. In speech, hesitation pauses are known to occur at points of high uncertainty (Goldman©Eisler, 1972). Clearly, making people conscious of their routine speech will slow down or interrupt the flow of speech, because previously parallel automatisms are now channeled through the limited™capacity bottle©neck; thus hesitation pauses may reflect high conscious involvement. There is considerable independent evidence for limited©capacity©loading events at junctures in the flow of speech, such as clause and sentence boundaries (Abrams & Bever, 1969). These junctures are likely to be points of high uncertainty. While this evidence does not prove conclusively that there is more conscious involvement at these points, it makes the hypothesis plausible.

Given these considerations, we can now explore the ideomotor approach to voluntary control.

7.3 The ideomotor theory in modern garb.

James' ideomotor theory fits neatly into the global-workspace framework. According to this view, a single conscious goal-image, if it does not meet with competition, may suffice to set off a complex, highly coordinated, largely unconscious action. For William James the ideomotor concept emerged from a puzzle in the experience of starting an action: Do we ever experience any command at all? Introspective reports on action commands were vague and contradictory, and this question became a‹j ‹ major source of controversy between Wundt, James, and the Wuô"rzburg School (James, 1890, pp.). James suggested that there is, in fact, no experience at all of commanding an action; rather, an action is organized and initiated unconsciously, whenever a certain goal image becomes conscious without effective competition.

We can partition the ideomotor theory into five interacting hypotheses:

1. The îConscious Goal Imageï Hypothesis is the one just stated, that all actions are initiated by relatively simple and momentary images of the goal. For many actions these images may be visual, because the visual system is very good in representing spatial properties of action. However, auditory, tactile, taste, or smell images are not ruled out. The act of walking to the kitchen to prepare lunch may be initiated by a taste and smell image of an attractive peanut-butter-and-jelly sandwich.

2. The îCompeting Element Hypothesisï is the notion that competing events may drive the goal image from consciousness. Competing events include conscious goal images as well as the non-qualitative intentions we have discussed previously (6.0). This idea has many important implications. It allows new conscious thoughts or images to interfere with the planning of an action, and it also permits editing of the goal by many different intentional goal systems.

3. The îExecutive Ignoranceï Hypothesis suggests that most detailed information processing is unconscious and that executive processes have no routine access to the details of effector control (Greene, 1971; Baars, 1980). Control of the muscles that are used to carry out an action is simply unconscious.

4. The îAction Fiat Hypothesisï claims that the moment of willingness to execute the action may be conscious, especially when the time to execute is non©routine. (James calls this the "fiat," the mental permission to start the act).

5. Finally, the îDefault Execution Hypothesisï is the tendency of the goal image to execute in the absence of any effective competition --- "by default." This is really just another side of the Competing Elements Hypothesis, but it is useful to emphasize it with a special label.

In addition to these five points, we should be reminded that îsubïgoals needed to accomplish the goal may become conscious if the goal cannot execute automatically (7.31). But let us suppose for now that all subgoals are automatic and unobstructed, so that they can execute without further conscious involvement.

To make these abstractions easier to imagine, take the example of retrieving a word, intending to say it, and then saying it. We have previously noted that complex activities like‹j ‹ word retrieval and speaking involve many separable components. Because of the limited capacity of consciousness we cannot afford to think consciously about many details in the act of speaking; we want to access all components of speaking at once, so that "the speech system" behaves as a single processor. But when we change from speaking to listening, or from speaking to eating, we may want to îdecomposeï the unitary speech system, to reorganize its components into new configurations for listening, chewing food, inner speech, and the like.

The ideomotor theory suggests that the "speech processor" as a whole must be recruited, organized, and triggered by a single conscious goal image. This image is itself controlled by a higher-level goal structure --- for example, the reader's general willingness to go along with demonstrations in this book. The following example explores ideomotor control in detail.

î1. Conscious Goal Images can activate unconscious goal structures.ï

If we ask the reader: "What are two names for the winged dynosaurs that lived millions of years ago?" The question is obviously conscious. Now, according to the ideomotor theory, this conscious experience initiates an intention to retrieve a word that matches the intention. Further, the conscious question triggers unconscious search processes which produce candidate words that may match or mismatch the intention (6.x). Because the words are rare, the momentary intention is likely to be prolonged into a tip-of- the-tongue state.

In GW theory, a "conscious goal image" is of course a global, consistent representation that provides information to numerous specialized processors (2.2). It is not surprising that a conscious goal would trigger local processors that control the muscles that carry out the goal. Indeed, we have argued early on (1.xx) that specialized processors are often goal-addressible: they are activated by goals. One nice feature of the GW system is that the goal image can be quite arbitrary or fragmentary, since it is the specialized processors themselves that have the real "intelligence" of the system, and which interpret the implications of the goal image in their own ways. Note that the goal image can trigger both the subordinate specialists able to carry out the action îandï the intentional goal context which constrains planning and execution without itself becoming conscious (Figure 7.3).

©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©
Insert Figure 7.3 about here.

©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©©

In fact, the goal image itself results from yet a higher-level goal context. Speaking is normally in the service of some‹j ‹ other goal --- communicating a thought, calling attention to oneself, gaining information --- which is, in its turn, in the service of even higher-level goals.

î2. Conscious Goal Images can also recruit a coherent set of action schemata and effectors, even though we have Executive Ignorance about the details of those processors.ï.

Now let us suppose that the reader has recalled the name "pterosaurus" (or "pterodactyl") as an answer to the question posed above. This is a conscious representation of the word. Now, how do we recruit the largely unconscious systems that control pronunciation of this difficult word? It is useful to recall here how complex and fast-moving the speech apparatus really is, and how little of it is accessible to awareness at any single time (Executive Ignorance). It seems plausible that the conscious word, in combination with a goal context, can recruit and organize the complex effector system needed to pronounce it.

îExecutive Ignorance of action detailsï is implicit in GW theory. As long as the details of action are unconscious, GW theory suggests that executive goal systems operating through the global workspace do not have direct access to such details.

î3. Default Execution: Given a compatible Dominant Goal Context, a conscious goal tends to execute automatically.ï

Once "pterosaurus" becomes conscious in the presence of an intention to say the matching word, something rather magical happens: we suddenly notice that our mouth has begun to pronounce the conscious word. The intervening steps of motor control are simply not conscious. In James' words, "consciousness is impulsive" --- unless, of course, other goal systems begin to compete for access to consciousness.

The notion that specialized processors tend to execute automatically, in the absence of contrary conscious messages, is already implicit in basic GW theory. There is nothing to stop an unconscious processor from executing an action except for contrary conscious images and intentions. If those are absent, we can expect actions to run off by themselves.

î4. The Competing Element Hypothesis: Conscious contents can be edited by multiple unconscious goal systems.ï

Suppose the reader first retrieves "tyrannosaurus", instead of "pterosaurus"? Clearly we do not want to execute this incorrect goal image. Various knowledge sources should interfere with its execution: some may remind us that "tyrannosaurus" is‹j ‹ too long, or that it has a different meaning. Such contradictory knowledge should have access to the global workspace, to compete against the incorrect conscious goal image. GW theory thus suggests that editing of flawed conscious plans is not some "added-on" capacity, but an integral aspect of the architecture of the cognitive system.

In GW terms, the goal image may also set off processors that generate competing goal images. Perhaps some of these contradict the first goal image, or present alternatives to it (see Figure xx). If some unconscious system detects a bad error in the goal image, it may trigger competing images that act to destroy the flawed conscious goal --- to edit and correct it. But once a single goal image wins out long enough, the dominant goal image will be executed. Its details are obviously off the global workspace, and hence unconscious. Figure 7.x presents this series of events in detail.

GW architecture supports editing of a global plan by potentially îanyï rule system. Take a single sentence, spoken by a normal speaker. Errors at any level of control can be detected îifï the sentence becomes conscious (e.g. MacKay, 1980). There are so many ways errors can creep into a sentence, and a correspondingly large number of unconscious rule systems that constrain successful sentences. There are many ways to be wrong, and only a few ways to be right by all criteria. Thus we can very quickly detect errors or anomalies in pronunciation, voice-quality, perceived location of the voice, acoustics, vocabulary, syllable stress, intonation, phonology, morphology, syntax, semantics, stylistics, discourse relations, conversational norms, communicative effectiveness, or pragmatic intentions of the speaker. Each of these aspects corresponds to very complex and highly developed rule-systems, which we as skilled speakers of the language have developed to a high level of proficiency (e.g. Clark & Clark, 1977). Yet as long as we are conscious of the spoken sentence we bring all these rule-systems to bear on the sentence --- we can automatically detect violations of any of them, implying that the sentence is somehow available to all of them (2.xx).

In principle, the set of "editing systems" is an open set. We can always add some new criteria for correct performance. This is one reason to suggest that conscious goals are "universally edited." Obviously the most effective competition is from goal contexts in the Dominant Goal Hierarchy, since these already have GW access during preparation and execution of the action (4.xx; 6.xx; Figure xx). But entirely novel aspects of the action can in principle be monitored and edited by onlooking processors, providing they can compete for access to the global workspace. Thus if one prepares to say a sentence, and suddenly a buzzing fly darts into one's mouth, the action can be aborted even though this situation was not anticipated as part of the goal context. Novel considerations can compete against the global goal.

‹j ‹å If all conscious goal images are inherently edited by onlooking processors, it follows that conscious goals that are actually carried out îmust have been tacitly editedï by relevant systems. Further, because îanyï system can potentially compete against the goal image, we can talk about this system as allowing îuniversalï editing. In section 7.32 we argue that this is indeed the criterial property of voluntary action: Voluntary action is action whose conscious components have been tacitly edited prior to execution.

î5. The Action Fiat Hypothesis: The moment of execution ââmay be under conscious and voluntary control.ï

We wait to say "pterosaurus" until we get a conscious signal; by contrast, in speaking a stream of words, we rarely seem to control the onset of each individual word consciously. But with an isolated word or action, given enough lead time, we can report fairly accurately our intention to execute the action at some specific moment. One key difference is whether the moment of onset of the action is automatically predictable; if it is, it is rarely conscious; but if the moment of onset is unpredictable, conscious control becomes more likely.

How should we represent the Action Fiat Hypothesis in GW theory? If goal images tend to execute automatically, it makes sense to suppose that timing an action involves inhibiting execution up to the right moment, and then releasing inhibition. Presumably, specialized processors sensitive to timing act to hold up execution of a goal image until the right moment (Figure x).

î6. Mismatch, surprise, and corrective feedback.ï

Conscious feedback resulting from an action can reveal success or failure to many unconscious goal systems, which may then develop corrective measures.

Imagine trying to say "pterosaurus" and actually saying, "ptero... ptero ... pterosaurus" --- a momentary stutter that is quite common in normal speech. Although we have no routine conscious access to the complex articulators and timing systems that control speech, it seems that those specialized systems îdoï have access to conscious events. In general, when we allow errors to become conscious, chances are that we can learn to avoid them in the future. In GW theory, consciousness of feedback from the flawed action sets into motion unconscious specialists that attempt to repair the dysfluency.

When we notice a speech error consciously, we often "repair"‹j
‹
it quickly (Clark & Clark, 1977; MacKay, 1980), but we are rarely conscious of details of the repair. Responding to overt errors is similar to anticipatory editing of covert errors, except that editing takes place before the action is executed (7.32). Correction of overt errors is useful in preparing for a more error©free performance next time around.

We have previously suggested that surprising events may involve disruptions of one level of context, even while higher levels are undisturbed (x.xx). Thus repair of contextual violations may start at a higher level than the level that was violated. The same thing may be true of errors in action. If we stutter, the error is at the level of articulation, but higher levels of control --- phonemic, lexical, syntactic, etc. --- are unaffected. Thus higher-level goal systems may seek another way to reach their goals. It is rarely the case that the entire dominant goal hierarchy is disrupted, fortunately for us (9.xx).

In sum, the Jamesian ideomotor theory can be incorporated straightforwardly in GW theory. In fact, it is difficult to see how one could believe that a conscious goal image is executed unconsciously without the concept of a distributed system of intelligent processors, able to interpret and carry out the relatively crude conscious goal.

The tip-of-the-tongue experience for "pterosaurus" helped to illustrate the intuitive plausibility of the ideomotor theory, and the rather nice fit with GW theory. But it does not provide proof. In section 7.4 below we will discuss the evidence for or against each hypothesis, and its implications for a broad theory of voluntary control. But first, we are ready now to make a basic theoretical claim about the nature of voluntary action.

‹F ‹

7.32 Voluntary action involves tacit editing of conscious goals.

If there is indeed universal editing of conscious goals, the conscious aspects of any action îmust have beenï tacitly edited for consistency with one's goal hierarchy before the action was performed. Take the example of premeditated murder. If a normal, rational person has thought for weeks about committing murder, and proceeds to do so, we immediately make the inference that contrary thoughts îmust have beenï entertained and rejected: that the murderer must have anticipated the chances of being caught, the likely disapproval of others, and perhaps the suffering of the victim and his family. That is, we immediately infer that competing alternatives will have been evaluated for any conscious impulse to action that was considered for some time, especially if the action has heavy potential costs. If the action was taken in spite of these "editing" thoughts, we make inferences about the value system of the murderer, or about mitigating circumstances. The important point for us here is the idea that conscious impulses are presumed to have been edited before action, assuming there was enough time to do so.

What components of action are likely to be conscious, and therefore tacitly edited? The theoretical answer fits our previous supposition (7.22): we tend to be conscious of those aspects of action planning that are novel, informative, significant, or conflictful (see chapters 4.0, 5.0, and 6.0). Those features that require the integrative capacities of a global workspace system are precisely those that are likely to be conscious ©©© and those are of course exactly the ones that are likely to need editing.

A major claim in this chapter is that îvoluntary action is in its essence, action whose conscious components have been edited before being carried outï. In contrast, counter©voluntary actions such as slips are editing failures: actions that îwould have been edited and changedï had there been an opportunity to do so before execution. Of course, most components of a normal action are unconscious; these components cannot be globally edited before execution. However, even automatic components of action must have been conscious at some time in the past. Therefore they must have been implicitly edited at that time to make them consistent with the Dominant Goal Hierarchy. Of course if automatic components run into trouble, they tend to become conscious, and can be edited.

There is direct evidence for an editing capability of this kind, and when we turn to the voluntary-involuntary contrasts

(7.xx) we will find that the major difference between closely matched voluntary and counter©voluntary actions is this ability to edit.

‹j ‹å The five main parts of the ideomotor hypothesis seem to fit the GW framework remarkably well. The interpretation does not seem forced or awkward. Further, as we look at the world from the resulting point of view, many other pieces of the puzzle begin to fall into place (7.x). The payoffs of bringing the ideomotor concept into our model are therefore attractive, and the theoretical costs seem minimal. But what about empirical support? Do the facts justify our taking the ideomotor theory seriously? Let us see.

7.4 Evidence bearing on the ideomotor theory.

A good deal of evidence is consistent with the ideomotor theory, though the case is not air©tight. Consider the following points:

7.41 Evidence for the impulsivity of conscious goal images.

The "Chevreul pendulum," a classic demonstration of the impulsive force of conscious goals, has been used has been used since the 19th century to persuade hypnotic subjects of the power of their own unaided thoughts (James, 1890). One simply takes a pendulum consisting of a string with a weighted bob at the end, and tries to hold it completely steady. Now, while trying to keep it steady, the subject begins to îthinkï of the pendulum as swinging away from and toward him, on a North-South axis. Without any perceived effort, the pendulum will begin to swing North and South. Again, making every effort not to move the pendulum, the subject begins to imagine it swinging right to left, in an East-West direction. The pendulum soon begins to follow the subject's thoughts, even though there is no noticeable effort or movement of the hand! It seems as if conscious images are more powerful than deliberate intentions.

It is not easy to adapt this classical demonstration to the rigors of modern investigation. The ideomotor theory needs a great deal more empirical support than is provided by demonstrations such as this. But it is difficult to doubt that there are conscious events related to goals: people can report their own conscious thoughts and images regarding a planned action, and usually predict their actions accurately in the very short term. But do those conscious events actually trigger off actions? This is difficult to be sure about, especially in view of the fact that some fleeting conscious goals that are difficult to report may evoke action (1.xx).

We do know that there is a momentary increase in mental workload immediately before the onset of an action (Keele, 1973). This is consistent with the idea that there is at least a momentary conscious goal prior to action. Libet (1985) has‹j ‹ presented arguments that we may become conscious of an action only îafterï the brain events that immediately trigger it. But this cannot be true in every case: surely there are many cases where people are conscious of what they are about to do seconds or hours before they do it, as shown by the fact that they can accurately predict their actions beforehand. The reader may make a conscious and reportable decision right now to turn the page, and actually do so: This is hardly surprising, but any theory that cannot handle this elementary fact is incomplete.

More evidence for the influence of conscious goals comes From the experimental literature on mental practice, showing th consciously imagining an action can improve performance as much as actual physical practice (Maltzman, 19xx). Conscious imaging of goals is used extensively in clinical practice and to improve athletic performance (Singer, 1984). There is no doubt that conscious images of goals can have powerful influence on effective action.

Further, we know that the opposite case also holds: îlossï of conscious access to an action can lead to a loss of control. Langer and Imber (x.xx, 1979) showed that automatization of a coding task leads to a loss in ability to evaluate one's own performance, and Reason's analysis of errors and accidents also shows a plausible relationship between automaticity and loss of control (x.xx, Reason, 1984). Automatization presumably means that goal images become less and less available, and therefore the actions themselves become less and less modifiable (x.xx).

Some of the most direct evidence for the role of conscious events in influencing action comes from conscious priming of experimentally evoked slips of speech and action. There are now several techniques for eliciting these slips in the laboratory (e.g. Baars, 1980, 1985, in press). One of these techniques uses phonological priming --- that is, conscious exposure to words that resemble the slip --- to elicit spoonerisms. Here is an example. The reader can ask someone to repeat the word "poke" about half a dozen times, and then ask, "What do you call the white of an egg?" Most people will answer, "the yolk" even when they know better. They have evidently been primed by the conscious word "poke" to retrieve a similar-sounding word from memory (Kimble & Perlmuter, 19xx). This technique may work because it duplicates the normal effect of conscious goal images, which prime the action to be taken.

In general, spoonerisms can be elicited by consciously priming the speaker with word-pairs that resemble the predicted error (Baars, 1980a, in press). Thus the slip îbarn doorï - îdarn boreï can be elicited by showing a subject a series of word-pairs like îdart boardï, îdark bowlï, îdot boneï, etc. Because subjects do not know ahead of time which word-pair they must say out loud, they must be prepared to say each one. This state of readiness apparently primes the system to make an error when the phoneme pattern is switched. ‹j ‹å

There are several other techniques for eliciting errors. All of them seem to create competing speech plans, compelling subjects to choose very quickly between the two alternatives (Baars, 1980b). Sentence errors like the following are triggered by creating uncertainty about the order of two phrases in a target sentence. If people are unsure about whether to say, îShe touched her nose and picked a flowerï, or îShe picked a flower and touched her noseï, they are likely to say inadvertently, îShe picked her nose ...ï. There are several ways to create this uncertainty. The easiest is to present the stimulus sentences, and after each one simply signal the subject either to repeat the previous sentence in the order given, or to reverse the phrases of the sentence. This technique produces predictable word™exchange slips at an adequate rate. Materials can be designed so as to elicit almost any involuntary statement from the subjects (Baars, 1980a; in press).

All slip techniques to date create a state of readiness in the speech system to act in a certain way ©©© they create goal contexts. Once this is done, we can ask whether adding a conscious image related to the target slip will increase the chances of the slip. For example, if we gave people the conscious word pair "terrible error," would that increase the chances of the slip "bad goof"? Motley & Baars (1979 a) showed that it does indeed. Further, if people are presented with a social situation such as the presence of an attractive member of the opposite sex, slips related to the situation are made much more often (see Baars, in press). In all these cases, a conscious prime coming just before a potential related slip will sharply increase the chances of making the slip. This suggests that conscious events can help recruit actions. While this evidence does not totally confirm the impulsive force of conscious goal images, it does support this part of the ideomotor theory.

7.42 Evidence for editing by global competition.

If a momentary conscious goal image is necessary to set up and trigger an action, competing conscious events should be able to delay or inhibit it. Everyday experience fits this pattern well. If we ask someone to remember a difficult word, and then interrupt with any other demanding conscious task, the desired word will simply not come to mind long enough to allow the person to say it. This is obvious and cannot be ignored. Thus editing may simply take place by competition for access to the global workspace, coming from processors that can detect the erroneous goal image. This competition can then keep the error from dominating the global workspace long enough to recruit and‹j ‹ trigger action. It is theoretically pleasing that we need add no new elements for editing to take place: it is simply another application of the general fact that the GW architecture permits local specialists to compete against global messages.

Other observations are consistent with this view. Thus Meichenbaum and Goodman (1971) have shown that impulsive children can use inner speech to improve self-control. If impulsivity consists of having very powerful conscious goal images that do not encounter immediate competition, then training children to use conscious inner speech may help them to compete against the undesirable goal image. The impulsive goal images may become less consciously available, and have less time to organize and execute unwanted actions. On the other side of the editing coin, Langer & Imber's findings (discussed above) indicate that practicing a task to the point of automaticity leads to a loss of ability to monitor the action. Apparently conscious goal images are less and less easy to monitor as an action becomes more and more automatic (Pani, 1982; see 1.xx).

Another source of evidence for anticipatory editing comes From experimentally elicited slips. One can get subjects to ma slips of the tongue that violate the general rules of language or usage; these slips can then be compared to very similar slips that do fit the rules. Thus, in the laboratory people will make slips like:

(1) îdarn boreï --- îbarn doorï (meaningful words)

(*) (2) îdart boardï --- îbart doardï (nonsense)

(3) înery viceï --- îvery niceï (syntactically correct)
(*) (4) îvice neryï --- înice veryï (wrong syntax)

(5) îlice negsï --- înice legsï (sexual comment that may be socially inappropriate)

(6) îreel fejektedï --- îfeel rejectedï (depressed comment)

Likewise, we can elicit word-exchange slips like:

(*) (7) îShe touched her nose and picked a flower.ï
--- îShe picked her nose ...ï (socially embarassing)

(*) (8) îShe hit the ball and saw her husband.ï
--- îShe hit her husband ...ï (aggressive affect)

(*) (9) îThe teacher told the myths and dismissed the stories.ï
--- îThe teacher dismissed the myths...ï (hard to pronounce).

(*) (10) îShe looked at the boy and talked softly.ï
--- îShe talked at the boy and looked softly.ï
(semantically anomalous).

(*) (11) îIs the gray sea below the blue sky?ï
--- îNo, the blue sky is below the gray sea.ï (false)

By designing slips that violate some level of control, and comparing them to very similar rule-îgovernedï slips, we have found a number of cases where the rate of rule-violating slips drops precipitously, sometimes even to zero (e.g., Baars, 1980a; Baars, Motley and MacKay, 1975; Baars & Mattson, 1982; Motley, Camden & Baars, 1979). All starred (*) slips listed above violate generic rules, and these slips show lower rates than similar slips that obey the rules. If the drop in rule-violating error rates is due to some editing process, the fact that this occurs with so many different rule-systems --- pronunciation, phonology, lexical, syntactic, social, etc. --- supports the idea of îuniversalï editing.

Ordinarily we think of "editing" as a review process in which someone like a newspaper editor checks the output of a journalist against certain criteria --- criteria like linguistic adequacy, fit with editorial policy, and the like. In general "editing" seems to involve two separate entities, one of which is able to detect errors in the output of the other system.

To show that editing in that sense occurs in normal speech production, we need to demonstrate that people in the act of speaking can detect mismatches between a speech plan and their criteria. Motley, Camden & Baars (1982) report that for a task eliciting sexually expressive slips (îlake muvï - îmake luvï, îbice noddyï - înice bodyï), there is a large and rapid rise in the electrical skin conductivity on sexual slip trials îeven if the slip is not actually madeï. On neutral control items there is no such effect. Since the Electro-Dermal Response is one of the standard measures of the Orienting Response --- a reliable physiological index of surprise --- these results suggest that a mismatch was detected even when the slip was successfully avoided. Thus egregious errors can be detected even before they are made overtly, and suppressed. This is exactly the notion of editing suggested above.

We cannot be sure in these experiments that the edited speech plan was conscious, but we do know that conscious speech errors can be detected by many largely unconscious criteria. Not all errors in spontaneous speech are detected, not even all overt errors (MacKay, 1980). But once speakers become conscious of an‹j ‹ error they are likely to correct it. In fact, normal speech is marked by great numbers of overt self-corrections or "repairs" (Clark & Clark, 1977). In any case, only part of the process of error-detection and correction is conscious and reportable. Certainly the slip itself is so, often, but detailed mechanisms of detection and correction are not. Therefore, even though we do not know for sure that the edited slips in the above experiments were conscious, we can certainly suggest that unconscious editing of conscious errors occurs quite commonly.

7.43 Evidence for Executive Ignorance.

Try wiggling a finger: where are the muscles located that control the finger? Most people believe that they are located in the hand, but in fact they are in the forearm, as one can tell simply by feeling the forearm while moving the fingers. What is the difference between pronouncing /ba/ and /pa/? Most people simply don't know. In fact, the difference is a minute lag between the opening of the lips and the beginning of vocal cord vibration. These examples can be multiplied indefinitely. We simply have no conscious, reportable access to the details of action.

7.44 Evidence for the Action Fiat.

We can prepare for an action and suspend execution until some "go" signal. The time of the "go" signal can be conscious, witness the fact that people can tell us when they will execute the action. In that sense, people clearly have conscious access to, and control of, the "action fiat".

The separation between îpreparationï and îexecutionï seems to exist even when execution is not delayed. All actions seem to have these two phases. For example, in the cat, where the neurophysiology of action control has been worked out to a considerable extent, there seems to be a natural division between preparation and execution. As Greene (1972) writes:

> "When a cat turns its head to look at a mouse, the

angles of tilting of its head and flexion and torsion of its neck will tune spinal motor centers in such a way that its brain has only to command 'Jump!' and the jump will be in the right direction. ... the tilt and neck flexion combine additively to determine the degrees of extension of the fore and hind limbs‹j ‹ appropriate to each act of climbing up or down, jumping onto a platform, standing on an incline, or peering into a mousehole; the neck torsion regulates the relative extensions of left and right legs when preparing to jump to the side. These postures must be set as the act begins; for if they were entirely dependent upon corrective feedback, the cat would have stumbled or missed the platform before the feedback could work. A few of these reflex patterns of feedforward are adequate for the approximate regulation of all feline postures and movements required in normal environments for a cat..."

When is the action fiat conscious? We can suggest that this depends on predictability of the time of action, just as consciousness or automaticity in general depends upon the predictability of any action subsystem. The action fiat should be conscious when the time of execution is non-routine.

7.45 Evidence for Default Execution.

How do we know that conscious goals tend to be executed in the absence of contrary conscious or intentional events? Part of the reason comes from the kind of demonstration of automaticity we suggested before: try looking at a word without reading it, or in the case of rapid memory scanning, try stopping automatic memory search before the end of the list. (Sternberg, 1963; x.xx)

Or consider once again the absent-minded errors collected by Reason and his colleagues (5.xx). Reason reports that îstrong habit intrusionsï occur in the course of normal actions when the actor is absent©minded or distracted, hence unable to pay attention (to be conscious of the relevant aspect of action). These cases argue for Default Execution. It seems as if a prepared action executes even when it should not, if contrary conscious events do not block the faulty action. This failure to block a faulty goal image can have catastropic consequences. Reason (1983) has analyzed a number of accidents like airplane crashes and road accidents, and concludes that many of these disasters may be caused by the intrusion of automatic processes, in combination with a low level of conscious monitoring.

A child of six knows how to keep such errors from happening: you have to îpay attentionï to what you're doing. That is, you must be conscious of the novel circumstances and goals. When we pay attention, erroneous Default Executions do not occur. However, it seems likely that the same principle of Default Execution is used to execute îcorrectï actions most of the time. We seem to automatically carry out conscious goals, unless contrary images and intentions block the conscious goals.

‹j ‹å

7.5 Explaining the voluntary-involuntary contrasts.

Earlier in this chapter we suggested that any complete theory of volitional control must explain the difference between the voluntary-involuntary contrasts: similar-seeming pairs of actions that differ only in that one is experienced as voluntary while the other is not (Table 7.1). Three categories of contrasting facts were explored in detail: the case of slips, of automaticity, and of psychopathology. Here we attempt to show how the theory we have developed so far can handle these facts.

An involuntary action that to escape voluntary control. It is often known to be wrong at the very moment it is carried out. We may hit a tennis ball with the sinking feeling that it is going awry, and yet our own psychological momentum may be unstoppable. Or we may make a slip of the tongue that never would be made if we only had a little more time to think (Dell, 1986;

Baars, in press). When we make an error out of ignorance or incapacity we do not speak of involuntary errors, errors that we know are errors, and that îîîïwould have been avoided îexcept forï --- what? One plausible explanation is that involuntary errors involve a failure of anticipatory editing, as described above. Editing occurs when systems that have spotted a flaw in a conscious goal begin to compete for global access to keep the goal from executing; but this editing function fails to work in the case of slips, unwanted automaticity, and the persistent errors of psychopathology. How could this happen?

Consider how editing might fail in our three primary cases: slips, automaticity, and psychopathology.

7.51 Slips: a losing horse-race between errors and editing.

If conscious goal images tend to be carried out by default when there are no competing elements, and if editing systems need time to compete effectively against faulty goal images, there must be a horse-race between "execution time" and "editing time." ("Execution time" can be defined as the time from the onset of the conscious goal image to the start of the action. Similarly, "editing time" is the time from the start of the goal image to the beginning of effective competition that stops execution of the act. See Figure 7.51.) In the case of slips, the editing systems lose the horse-race, because execution time is faster than editing time. The faulty action executes before editorial systems have a chance to compete against its goal image.

There is one obvious case where this may happen: we know that practiced images fade from consciousness or become very fleeting, and that highly practiced, predictable actions become more efficient and less conscious. Pani's (1982) studies on the automatization of images show this pattern. As we have discussed before (1.aa), Pani showed that conscious access to visual images used in solving a problem drops consistently with practice. In terms of our model, we can suppose that images become globally available for shorter and shorter periods of time, until finally they are globally available so briefly that they îcan no longer be reported, even though they continue to trigger highly-prepared effector systemsï. Highly prepared processors presumably can react very quickly, while the act of reporting goal images may take more time. Alternatively, it is possible that goal images are simply lost from the global workspace; that they are not even fleetingly available. In the remainder of this discussion I will assume that the first case is true --- that with practice, goal images are still broadcast globally, but more and more fleetingly. Naturally this hypothesis must be tested (see 7.xx).

If goal images become more and more briefly available with practice, the previously discussed studies by Langer and Imber (1979) begin to make sense. These authors found that more practiced subjects in a coding task were îmoreï willing to accept an incorrect assessment of their own performance than less practiced subjects. These authors argue that overlearning a task can reduce the knowledge the subject has about îhowï the task is performed, and under these circumstances subjects should be more vulnerable to negative assessments of their own performance, because they cannot evaluate their performance anymore by themselves. This is exactly what we would expect, given the assumption above that the goal image becomes less and less available with practice. Automatic, highly prepared effector systems can continue to carry out the task (because they became more well-prepared and efficient with practice, and therefore needed less of a goal image to be triggered). But asking someone to do something novel such as evaluating their own performance, should become more difficult, because the global goal image upon with this evaluation can operate is available only fleetingly.

Thus the goal image controlling the "counter©voluntary" act may be available long enough to trigger a îpreparedï action, but not long enough to be vulnerable to interference from editing systems.

In Figure 7.xx we show our usual model, with the goal image G able to trigger off processors that tend to carry out goal G, barring competing messages from other systems that may not approve of G, which we will call ~G ("not-G") messages. If G is globally available only very fleetingly, but long enough to trigger well-prepared processors, then editing may fail because the effectors may be faster than the editing systems. Systems like ~G in Figure 7.xx find it difficult to interrupt, and modify the goal image G. In this way, an action may "slip out" in an‹j ‹ uncontrolled way, because competing processors could not catch it in time. Goal image G can come and go very rapidly, because there are automatic systems able to execute it, and competing ~G messages are too slow to stop its execution.

Notice a very significant point here: there is a trade-off between îcompeting against Gï and îrepairing Gï. In order to correct G, to modify it, to suggest alternatives, and the like, it is important for many processors to have global access to it. G must be available for a fairly long time if it is to be modified by other systems. This is of course the whole point of making something conscious, that many different unconscious experts can cooperatively work on it (2.xx). But ~G systems compete against G in order to stop its execution, and therefore make it îlessï consciously available. If it is less available, there is less time to modify G, and to improve it. This trade-off will be very important in our discussion of psychopathology below.

Figure 7.xx tells why faulty goal images may be carried out in spite of the fact that their faultiness is known; but it does not tell us why the inner error occurred in the first place. In the case of slips, I have argued in related work that competing goals are often the cause of errors (Baars, 1980; 1985). For example, there are often two different ways to express a single thought. The two alternative goal images may compete for global access, they may fuse or alternate. When there is a limited time to resolve this goal competition, errors are likely to occur, especially if other events load limited capacity at the same time (Baars, in press, d; Chen & Baars, in press; Dell, 1986).

The horse-race between execution time and editing time is key to the view of involuntary action we will maintain in this discussion. It has strong implications not only for understanding slips of the tongue, but also for unwanted automaticity and psychopathology.

7.52 Counter-voluntary automaticity: the case of "structural" slips.

Once the triggering conditions for any automatic process are provided it becomes difficult to stop voluntarily. Habitual cigarette smoking has an involuntary quality, as do compulsive eating, nervous movements, and the like. Once we simply look at a word, it is essentially impossible to stop reading it. The large experimental literature on these phenomena makes the same point (Shiffrin & Schneider, 1977; LaBerge, 1984). Habit is indeed the "great flywheel of society," in James' well©known phrase, and there are times when the flywheel runs out of control, resulting in fatal accidents (Reason, 1984). Whenever we try to resist an automatic habit, it will start to behave "erroneously" with‹j ‹ respect to our purpose. Such "errors" have much in common with the slips discussed above. Just as in the case of slips, automatic execution time is plausibly faster than voluntary editing time. Thus we can apply the same "losing horse©race" model in the case of unwanted automatisms. They seem to reflect the same mechanism.

Of course, automatisms are not immune to change. Changing them often requires repeated efforts. It may often be helpful to block or slow down the action to make it more conscious and easier to edit and modify. To speed up the editing process, we may we need repeated trials to improve automatic editing and re-establish voluntary control.

7.53 Psychopathology: the case of repeated editing failure, perhaps due to excessive control effort.

If we are going in the right direction in this discussion, what can we say about repeated errors that are îknownï to be errors --- the case of psychopathology? The voluntary system we have explored so far aims above all to achieve goals and minimize errors. But in psychopathology we find a great range of behaviors that violate voluntary goals, and that repeat errors with remarkable persistence. Psychopathology seems to involve an enduring failure of the entire voluntary control system. How could such persistent failures arise?

Above, the lack of conscious availability was thought to be responsible for a loss of editing ability. We block and fix errors by making them available longer. If there is a repeated editing failure in pathological symptoms, what could stand in the way of this normal editing process?

One answer may be that îthe very attempt to block wrong goal images may stand in the way of adaptation to the errorï. We have referred above to the trade-off between modifying a goal îGï and blocking its execution. That is, if we block a goal image, we stop the goal from executing, but we also lose the opportunity to modify and improve it. In order to fix a faulty goal image, we must be allow it to be conscious for some time. But in the case of pathological errors, editing systems may attempt to wipe the goal image from consciousness as quickly as possible. In psychopathology we may be trying to block the faulty goal image so quickly and completely that we have no time to fix the problem.

Take the example of a fearful image of an airplane crash. Every time we think about taking an airplane trip, we may have a vivid fearful image of the plane going down in flames. If we allow ourselves to contemplate the image for a while, we may notice that we can also mentally reverse the plane crash ©©© its flaming wreckage may turn imaginatively into a whole new‹j ‹ airplane, and leap back up into the sky to continue its journey. Just by allowing the image to remain conscious, many unconscious processors will have access to the image. These unconscious processors may be able to modify the conscious image in various ways, thus creating a greater sense of control (Singer, 1984). The problem may come when we do not allow ourselves to contemplate the fearful image at leisure. Rather, we edit it quickly so as not to deal with its awfulness (Beck, 1976; Ellis, 1962). In that case, we do not provide the time needed to change the image, to create alternatives, and the like. Then the fearful mental image may become a rapid, frightening, and uncontrollable phobic thought. It is this trade-off between "editing by competition" and "fixing by conscious exposure" that may cause phobic images to take on a life of their own.

If that is true, then allowing the phobic image to become fully conscious, changing it to a safer image, and in general, gaining more voluntary control over it --- all these methods should work in the control of phobia. And indeed, these techniques are the essence of phobic control: systematic desensitization, imagery techniques, flooding, and practicing the symptom may all work by allowing the phobic image to remain conscious long enough to notice that the reality is not as awful as the anticipation.

From this point of view the "paradoxical" techniques that are sometimes so effective take on great importance. Voluntarily getting children to stutter apparently solves the problem in some cases; asking phobics to practice fearful imagery may help that problem, and so on. These results make perfect sense from our perspective: voluntary stuttering presumably causes a goal image to remain conscious for a longer time, without destructive competition to reduce its duration. And if it is available longer, other systems can act upon the goal image to modify it, so that it comes under the control of systems which failed to control it before. Paradoxical practice of the to-be-avoided action increases our ability to avoid the action.

It would be foolhardy to claim that this is the only mechanism of psychopathology. But it may be one central factor that sustains and aggravates a variety of repetitive dysfunctional behaviors. It has great simplicity, there is some good evidence for it, it is quite testable, and it flows naturally from our entire discussion in this chapter.

In summary, we have explained the contrastive facts shown in Table 7.1 by means of a modern ideomotor theory. It is plausible that voluntary control is guided by momentary goal images, even though those images are difficult to assess directly. The five major points of the ideomotor theory seem to have some empirical support, though more is needed. There is a satisfying fit between the ideomotor theory and the theoretical approach to consciousness we have been pursuing throughout this book. As we see next, the ideomotor theory seems to generate fruitful‹j ‹ hypotheses about a number of problems, including the nature of decision©making, perceived effort and control, the nature of non™qualitative conscious contents, and even the understanding of absorbed states of mind and hypnosis.

7.6 Wider implications.

7.61 What does it mean to make a decision?

Most of our actions are bound by past decisions that are not currently conscious. As children we learned to pronounce the difficult phoneme cluster /ask/ as "ask" rather than "aks," with a lot of conscious concern for the different sounds. Once learned, the difficulty fades into the background, and we need not make the same decision again. All actions contain the residue of commitments made at previous conscious choice-points, decisions that are no longer conscious. If the goal hierarchy has an established commitment to a certain decision, there is no need to become conscious of the excluded alternatives. On those potential choices we now have established policies.

But perhaps some aspect of almost any action is consciously decided --- its timing, its propriety in a particular situation, etc. Much of the time people can make voluntary decisions about consciously entertained choices. We can decide to read a chapter in this book, to adopt certain life-choices in adolescence, and occasionally we can even make clear and effective decisions to stop or start long-term habits. These are all choices with conscious alternatives. If consciousness is the domain of competition between such alternative goals, our model should be able to show how we make decisions that stick, as well as those that do not last.

The simplest approach is to say that one can broadcast alternative goals, like "Should I îGï ... ?" followed by "Or shouldn't I îGï ...?" and allow a coalition of systems to build up in support of either alternative, as if they were voting one way or another (Figure 7.61). The stronger coalition presumably supports a goal image that excludes effective competition, and which therefore gains ideomotor control over the action (7.0). Thus voluntary actions may be preceded by a long set of problem-solving triads, as described in Chapter 6.

Insert Figure 7.61 about here.

But where does the conscious goal image come from in the first place? If the goal hierarchy is not strongly violated, it‹j ‹ presumably does not generate conscious goals (4.0). In that case the hierarchy may îconstrainï the goals that are to become conscious, without producing its own conscious events. Sometimes of course the goal hierarchy is deeply challenged, and must generate conscious goal images to maintain its integrity, or to prepare for change (x.xx). Further, some conscious choices are presented by the outside world, as when someone offers us a tempting dessert, an attractive item on sale, or a career opportunity. Other conscious choices are surely created by internal changes, like the beginning of hunger or the onset of puberty. Some may be created by unresolved conflicts between deep goal structures, like the need to control others vs. a desire to be liked by them. And some conscious choices may be generated by a continuous process of entertaining long-term dilemmas that have no simple solution.

All these points raise the issue of îindecisivenessï. As James knew so well, the question of getting out of bed on a cold morning appears as a struggle between alternatives. Perhaps most of our ordinary decisions have this quality. But some extended struggles may be won by patience rather than force. As James noted in the epigraph, one can simply wait until the cold of the morning fades from consciousness; if a fortuitous thought about getting up then emerges, it may be able to dominate consciousness without competition. Thus the persistently returning thought ultimately wins out. The idea that important goal systems can "win out" by sheer persistence, by returning to consciousness again and again, is consistent with evidence from thought monitoring (Klinger, 1971; Pope & Singer, 1978), showing that unresolved issues tend to recur spontaneously.

Indecisiveness may be the case where neither of the two contending goals ever completely fades away. Young children often seem indecisive compared to adults. They may be quite impulse-driven, sometimes hesitating back and forth, back and forth, between two attractive goals. Young children may not yet have a consistent dominant goal hierarchy. Over time, many repeated cases of effective coalition-building between the most successful goals may result in the relatively stable adult goal-hierarchy, so that consistent goal contexts become established and automatized like any other skill.

We have noted that conscious goals that are consistent with the goal hierarchy will last longer than those that violate it

(7.83). It also follows from our current model that some conscious goal images may fit the deeper levels of the goal hierarchy better than the more superficial levels. In that case the more deeply- driven goal images may last longer, or they may return again and again until they lead to action. One way to make new goals effective is to tie them in with existing deep goals. Thus one may have an inner argument of the form: does my commitment to survival make it necessary to go out and jog four miles a day? Does my commitment to social success make it imperative to stay at this boring party? In these cases a‹j ‹ conscious connection is created between an immediate goal and an existing deep commitment; in just this way politicians will make a case for new expenditures by reference to existing deep and agreed-upon goals like "national security," "winning the war on poverty," and "bringing back prosperity." By consciously mobilizing the deep goal hierarchy, one's superficial reluctance may be overcome. These rhetorical connections between local goals and deep goals may be specious, but as long as they allow the conscious goal to be available long enough to be executed, they will be effective.

Below we argue that most normal action is relatively conflict-free (section 7.83); that is, it takes place in the domain of overlap between many deep goals. But conflicting goals are not unusual. Any new major goal must of course be reconciled with existing priorities.

Thus much of the time people may carry on an inner argument about their goals. Not all of this inner argument may be fully conscious; some of it may consist of fleeting images that function merely as reminders. Notice an interesting thing about this inner argument: If the ideomotor theory is true, it is very important to îhave the last wordï in a train of arguments; the last word, after all, is the one that will be carried out, because it is not followed by competition.

Inner arguments about goals have many implications. For example, one can define a "belief" as an abstract concept that is not disputed in the stream of thought, though it could be. One can dispute a political or religious belief, or a scientific position, but it is quite remarkable how rarely people challenge their own beliefs. A belief îsystemï may be defined as a consistent set of such undisputed concepts, one that serves to stabilize one's thoughts, feelings, and actions --- presumably just by giving the belief system the last word in the inner argument, since the last word has the real power by the ideomotor theory --- it is the one that controls action without contradiction. Likewise, a îclosedï belief system is one that has a ready answer to all apparent counter-arguments, so that any possibility of change is minimized (Adorno, et al, 19xx; Rokeach, 1960). Perhaps all ideological, political, philosophical, and even scientific belief systems are closed to some extent. Simple observation should convince us that most people have self-serving closed belief systems about many disputable topics, especially those that are difficult to decide on direct evidence.

‹` ‹

7.62 Resistance to intended actions, perceived effort, and perceived voluntary control.

We noted above that voluntary control is different from spontaneous problem solving (6.0) in that we usually îknowï that our voluntary actions are goal-directed. We have "metacognitive" access to many voluntary goals, and often to mental events that block voluntary goals (e.g., Flavell & Wellman, 1977). Two factors may give us this kind of metacognitive access. The first is obvious: îex hypothesiï, the ideomotor theory states that voluntary control involves conscious images, which, if they are available long enough, are also available to metacognitive processors. Metacognitive processors are presumably involved in representing, recalling, and describing the fact that we do have a certain conscious goal. Thus the reader knows that s/he is reading this book voluntarily, in part because he or she may be able to recall the conscious goal of doing that.

But there must be many times when we experience an action as voluntary even when we do not remember its controlling goal image. After all, goal images are fleeting, their memory may be masked by later events, and so on. By our discussion above, the more automatic the action, the less we can report our intention. Conversely, the more the action îencounters resistanceï, the less automaticity will operate, and the more a decision to act can typically be reported.

This suggests that resistance to performing an action, the perception of effort, and perceived voluntary control are all of a piece. Take the issue of perceived effort, which appears very much as a conflict between expected control and actual control. The author's experience in typing the manuscript of this book may illustrate the point. As a practiced (though errorful) typist, I am normally unconscious of the details of typing. The computer program that displays words on my screen works so quickly that normally I do not notice it at all. But sometimes when I am typing, the computer is simultaneously printing out some other material, and then the screen seems to slow down enormously. The lag time between a finger stroke and a character appearing on the screen is now very long, compared to my expectations. Now the relationship between key-strokes and characters on the screen becomes agonizingly conscious. The subjective experience is one of great effort, as if I must forcibly push each character onto the screen. I am acutely aware of the voluntary character of every keystroke.

This example may provide a way to test the hypothesis that perceived voluntary control results from perceived effort. It suggests that a goal context contains information about the length of time an action should take. When this time is delayed, we tend once more to become conscious of both goal images and feedback, so that many processors can now operate on the conscious components of the action. As the conscious goal image‹j ‹ becomes more available, metacognitive processors can also operate on it, to facilitate recall and self-description. In this way, our knowledge that we have a certain goal may depend on violations of automaticity in accomplishing the goal.

Notice that the increase in conscious access to such a delayed goal gives us four distinct advantages: first, we have more time to edit and change the conscious goal; second, in this process we improve our voluntary control over the action; third, we can comment on the goal in question metacognitively, which then allows us to recall it, to talk about it, and perhaps to find alternative ways to accomplish the same ultimate end. Finally, as we will see in Chapter 8, access to a conscious goal can also guide îlaterï conscious contents, as when we make conscious decisions about what to pay attention to next.

7.63 Ideomotor control of conceptual thought: A solution to the puzzle of non-qualitative consciousness?

One of our persistent thorny problems has been the relationship between clear, qualitative conscious contents like percepts, feelings and images îversusï non-qualitative conscious events like concepts, beliefs, expectations, and intentions, which surely take up limited capacity, but are not experienced with qualities like warmth, color, taste, and smell (x.xx).

We have previously remarked on the fact that human beings have a great tendency to concretize abstract ideas: to think in terms of metaphors that can be visualized, or to reduce an abstract class of events to a concrete prototype. It may be no accident that mathematics and physics really have two separate symbol systems: an algebraic code and a geometric one. The two are mathematically equivalent, but not psychologically, because people can use their visual imagination with geometric graphs but not with algebraic formulas.

But we do not have to resort to science for examples. All of us clearly represent the meaning of a sentence in an abstract form. To illustrate this, let the reader recall word-for-word the sentence before this one. (No looking!). The great majority of readers will not be able to do this, but they will be able to recall a îparaphraseï of the sentence --- that is to say, a semantic equivalent, with different words, different syntax, and even different sensory qualities than the original sentence; but the paraphrase will preserve the abstract meaning of the original. The evidence is very good that educated adults rapidly convert words and sentences into a semantic code that is quite abstract and impossible to experience qualitatively (Bransford & Franks, 1976). The question we must face here is, of course: How‹j ‹ do we then manipulate the abstract semantic code through consciousness?

One possibility is that we have ideomotor control over abstractions. Take the complex conceptual structures developed in the course of this book. We have now defined îtermsï like "context," "global workspace," and the like, which are perceptual in nature even though they refer to abstract non-qualitative things. The more we can manipulate these abstractions with words we can see and hear, the easier we will find it to understand the theory. Likewise, we have made a great effort in this book to present useful metaphors for our theory, such as the image of a conference of experts, each competing for access to a central blackboard. But every concrete metaphor is really inadequate. The conference image fails to show that expert processors in GW theory are decomposable, while human experts are not (2.xx). (Human experts have a tendency to stop running when they are decomposed.) This point is quite general: metaphors are inadequate representations of a more abstract and complex reality (Lakoff & Johnson, 1980). In science, they must be used with care.

In general, an imageable metaphor seems to serve the function of evoking and recruiting conceptual processes that are more abstract and often more accurate than the image itself. These abstract entities may be impossible to experience qualitatively. Hence the need for visual figures, audible words, and concrete metaphors. These can be made qualitatively conscious when needed, to stand for abstract non-qualitative entities.

All this suggests that we do indeed have ideomotor control over abstract concepts, so that we can always concretize an abstraction, and conversely, we can always abstract from concrete symbols. It is not enough merely to translate the perceptual world into abstractions as we do in comprehending a sentence; in turn, we must be able to retrieve the abstractions in perceptual form in order to work with them, to resolve conflicts between them, to make predictions from them, and to use them to act on the world. In all these transformations, it is useful to recode the abstractions into some qualitative, imageable form. The ideomotor theory seems to add some real clarity to the problematic relationship between qualitative experience and abstract representation.

7.64 Fleeting goal images make accurate source attribution difficult.

If it takes time for a goal image to result in action, then what about goal images that not only trigger an action, but also require us to talk about them? That is, images that recruit îtwoï actions? If we want people to report their own goal images, they‹j ‹ must make the goal image available long enough not only to trigger the original action, but also to help recruit linguistic systems able to describe the image. This is of course the same problem we encountered before, when we first raised the possibility of fleeting conscious events that pass too fast to describe (1.xx). The best example, again, is in tip-of-the-tongue states when people experience the missing word fleetingly, and encounter the frustration of trying to hold on to the image long enough to say it.

This is a fundamental problem in metacognition. It may help to explain a number of problems in self-observation. There is of course the social©psychological literature on errors in attribution of personal causation, and the frequent failure of people to know their own reasons for doing things (Nisbett & Wilson, 1977; Weiner, 19xx). Accurate source attribution is very important for metacognitive knowledge and self-control. Again, this is not a topic we can explore in much detail, but we can suggest its relevance to GW theory.

7.65 The relationship between a goal image and the action it evokes may be highly variable.

The act of walking to the kitchen may be triggered by many different goal images. One can imagine a seductive peanut-butter-and-jelly sandwich, or left-overs from last night's dinner; one can remember that the stove needs cleaning, or imagine the odor of cooking gas. We need not imagine any of these in great detail. A fragment of a related image will do quite nicely to trigger a habitual action. This is very much like the issue of synonymy and paraphrase in language: there are dozens of ways of saying the same thing. In action control, a conscious cue is presumably interpreted by many different context-sensitive systems. We do not need a detailed conscious plan or command, since the action is carried out by specialists that know more about local conditions than we do consciously. Various unconscious specialists keep continuous track of our posture, balance and gravity, about salivation and digestive enzymes to prepare for eating, about remembering the route to the kitchen. Greene (1972) has pointed to the ambiguity of commands in distributed control systems as a general and very useful property.

This point has important implications for research. We must not fall into the trap of looking for îtheï goal image for walking, or talking, or for any other action that looks the same in different circumstances. This is what misled introspectionists like Titchener and Kuô"lpe, who were astonished to find the great range of variation in mental images between different observers‹j ‹ (1.xx). The modern ideomotor theory indicates that many different goal images can serve to recruit and initiate any given action. Conscious images may seem quite irrelevant, and still result in appropriate action. Imagining a sandwich while walking in the desert must not trigger an automatic walk to the kitchen, but it can stimulate new efforts to find food and water. Thus goal images may vary tremendously between different situations and observers, and yet be quite effective in controlling normal voluntary action.

7.7 Absorption and hypnosis as ideomotor events.

7.71 Absorption as a drop in competition for GW access.

The ideomotor theory has many interesting implications. For example, it suggests a reasonable account of hypnosis as a state in which ideomotor control operates without effective competition (7.67). Before we discuss this, we can define îan absorbed stateï --- watching a movie, reading a novel, and the like --- as a state in which only one coherent stream of events dominates consciousness (viz., Tellegen & Atkinson, 1974; Spiegel, 1984). That is, there is a low level of effective competition between different topics (dominant contexts), and there is no voluntary effort to change topics (see Chapter 8).

In principle, it would seem that there are two ways to reach an absorbed state. One is for the number of competing contexts to decrease. This may happen simply when we relax, let go of our current concerns, solve a major preoccupying problem, or enter a state of trust that things will work out without voluntary effort (Klinger, 1971). A second way to enter an aborbed state is to allow one context to become extremely dominant and thereby to exclude alternatives. Shadowing tasks compel one to repeat immediately each word in a coherent stream of words (Cherry, 1953; Broadbent, 1958). This task is so demanding that competing thoughts are simply excluded from conscious experience. Nevertheless, competing thoughts have probably not disappeared; they are simply excluded from consciousness. Thus we can enter an absorbed state either if consciousness is dominated by a very strong context, or if there is a drop in competition from alternative contexts. In fact, of course, most actual absorbed states have both of these features. In watching a fascinating movie our experience is being structured by the story line, which continually generates new expectations about future events that need to be tested. At the same time we may relax, postpone some pressing concerns, and thus lower the urgency of competing topics.

One implication is that îwe are always in an absorbed state relative to our own dominant contextï. If we look at the goal hierarchy (x.xx), we can see that its lower levels can change‹j ‹ much more easily than higher goals, which are quite stable over time. Most people do not cease wanting to survive, to be socially accepted and respected, and to pursue other life-long goals. Adults change their major beliefs and goals quite slowly if at all. Even perceptual and imaginal contexts change only rarely. This suggests that we are never absolutely "absent-minded," "mindless," or even "preoccupied" (Reason & Mycielska, 1984). We are always "present-minded" to our îownï dominant preoccupations. Now, if we are driving a car and thinking thoughts of love at the same time, we may run over a pedestrian. îRelative toï the act of driving the car we were preoccupied and absent-minded. But relative to thinking thoughts of love, we were quite present. Taking a bird's eye view of the situation, it would seem impossible to be utterly absent-minded. "Absorption" is only a relative term.

When we are absorbed in one mental topic to the exclusion of others, the other topics must go on automatic. Thus if we were to ask someone to shadow speech while performing a fairly routine task --- driving a car along a familiar route --- we would see the automatic components of driving emerge with minimal conscious and voluntary overlay. We should then expect to find large numbers of automatic "habit intrusions" into the act of driving (Reason, 1983). Driving a car distractedly may be rather suicidal, but similar experiments can be done under less dangerous circumstances.

7.72 Hypnosis as ideomotor control without competition.

Absorption has long been thought to be a key element in hypnosis. When we combine the idea of absorption with ideomotor control, we have a possible theory of hypnosis (James, 1890, ch. on hypnosis). The major features of hypnosis seem to flow from the fact that in this state we have only one conscious goal image at a time, which tends to be carried out because the chances of competition from other elements are reduced. Although we cannot go into hypnosis in great detail, this possibility is worth exploring briefly.

What are the major features of hypnosis? There seems to be good agreement on the following:

1. îAbsorptionï, sometimes called "monoideism," or "imaginative involvement" (Tellegen & Atkinson, 1974; Spiegel & Spiegel, 1978; J. Hilgard, 1979; E. Hilgard, 1977; Ellenberger, 1970). Hypnosis seems to create a new, imaginative context that dominates experience for some time to the exclusion of other events (Singer, 1984).

‹j ‹å 2. îDissociationï. Good hypnotic subjects show several kinds
of spontaneous dissociation. First, there are two kinds of temporal dissociation. A good subject is often spontaneously amnesic for the experience, which is a kind of post-hypnotic temporal dissociation. There is also pre-hypnotic dissociation, since separation from previously dominant trains of thought is

common (J. Singer, personal comm.). In addition to temporal dissociation, two kinds of îconcurrentï dissociation occur. These may be called "dissociation from effectors" and "dissociation From the normal self." Subjects often report feelings alienation from their own limbs that are manipulated by suggestion, as if their arms and legs had "a will of their own" (Spiegel & Spiegel, 1978). Further, there is commonly some surprise at îoneselfï for allowing the hypnotic actions and experiences to happen, so that there is a kind of dissociation between one's "normal self" and one's "hypnotic self" (viz., 9.x). In sum, there is an experienced îtemporal separationï from earlier and later states, and also îconcurrent separationï during hypnosis from one's own normal experience of self and one's own hypnotically controlled actions.

3. îSuggestibilityï is a defining feature of hypnosis; but this can be viewed as a consequence of ideomotor control, plus a kind of dissociation from normal inhibitions. But dissociation From inhibiting thoughts is of course a property of absorptio as a state of low competition for consciousness. If the ideomotor theory is true, and if our normal editing mechanisms are not competing against novel conscious contents, it follows that one will show a great flexibility in obeying the conscious ideas. Thus suggestibility seems to flow from "absorption plus ideomotor control."

4. Strong and stable îindividual differencesï. About a quarter of the population is highly hypnotizable. These people easily slip into this remarkable state with a very simple, standard induction (Hilgard, 1977; Spiegel and Spiegel, 1978).

5. Hypnotic îinductions are arbitraryï: anything believed by the subject to induce hypnosis will induce hypnosis (James, 1890). However, relaxation and a reasonable feeling of trust for the hypnotist are common features.

These are the positive facts about hypnosis that command a good consensus. At the same time hypnosis has some îpuzzling "negative" featuresï: properties that we might expect, but which researchers have not found in spite of repeated efforts.

1. No reliable neural correlates of hypnosis have been found so far. Physiologically, hypnosis looks like a normal waking state.

2. There is a puzzling absence of strong and reliable personality correlates, in spite of many attempts to find them (e.g., J. Hilgard, 1979). ‹j ‹å

3. It is difficult to demonstrate a conclusive difference between hypnosis and "pretense" (T.X. Barber, 19xx). But this may be in good part because very good "pretenders" are able to experience their pretended states very deeply and realistically --- in other words, good pretenders may enter an absorbed state, in which only one train of conscious contents dominates their actions. There may thus be no real difference between very good play-acting and hypnosis, but this fact may reveal as much about acting as it does about hypnosis. Many professional actors experience deep absorption and identification with the characters they play. The difference between half-hearted acting and Stanislawskian "method acting" is the difference between being superficially involved, and being deeply absorbed in a certain character (Stanislawsky, 19xx). Stanislawskian method actors may believe for a while that they are the character they are playing. Absorption may be the key both to good hypnosis, and to good pretense as well.

îHypnosis as absorbed ideomotor control.ï

Several investigators maintain that absorption may be the basic element of hypnosis, the single central fact from which all else flows (e.g. Spiegel and his refs). This is a very attractive argument from our point of view. We can simply take James' ideomotor theory, in our modern version, and explain all of the features listed above. That is:

1. îAbsorptionï or monoideism is simply a low level of competition for access to consciousness between alternative contexts. Under these conditions of "low editing" the dominant stream of consciousness, which may be quite different from our normal states, should be in control. *** ALSO metacognition, see next Chapter.

2. îDissociation.ï We can easily explain both temporal and concurrent dissociation. Spontaneous amnesia after hypnosis is a difficulty in voluntarily reconstructing the hypnotic state in such a way as to easily retrieve information from it. This is not surprising, given the differences in content between the hypnotic, absorbed context and our more usual contexts, those which we îcallï normal and not hypnotic. So spontaneous amnesia would seem to follow quite easily. The same argument applies to the temporal dissociation from the dominant context before hypnosis. We should feel dissociated from it, given the differences in content. If hypnosis is mainly an absorbed state, there should be relatively few shared features between it and our normal state, thus making recall difficult.

The two kinds of concurrent dissociation also make sense. Dissociation from our own hypnotically controlled limbs may be‹j ‹ just the act of noticing the truth of ideomotor control. In our normal waking state we tend to forget that are îalwaysï unconscious in detail of actions that carry out our conscious goals. That is what the ideomotor theory is all about, after all. Similarly, we are normally unconscious of the controlling contexts of our own actions. In hypnosis we may be surprised to realize that. But in our whole approach in this chapter, the unconsciousness of goal contexts and automatic actions has become a fundamental assumption about normal action. From this point of view, it is not dissociation that is unusual. What is novel in hypnosis is the fact that we îrealizeï the existence of ideomotor dissociation between conscious events, their consequent actions, and their antecedent goal contexts. Perhaps we recognize this in hypnosis because hypnotically controlled actions are often unexpected;

they violate our usual contextual assumptions about ourselves.

The close connection between absorption and dissociation comes out in Spiegel's (1984) clinical observation that

"... it has been commonly observed that many highly hypnotizable performers, such as actresses and musicians, dissociate their ordinary awareness of themselves when they are performing, and feel strangely disconnected from the results of their peformance after it is over. One highly hypnotizable pianist reported that her only memory of her graduation recital was of floating above the piano admiring the grain of the wood. She had to ask a stagehand whether she had in fact completed her program, which she had performed extremely well."

Other features of hypnosis are also consistent with this point of view.

3. îSuggestibilityï and flexibility seem to be merely the result of ideomotor control in an absorbed state, with minimal competition and self-examination.

4. The reasons for the strong and stable îindividual differencesï in hypnotizability are not clear. Given that the notion of hypnosis as "absorbed ideomotor control" seems to work quite well, the question seems worth exploring from this theoretical perspective.

5. The îarbitrarinessï of hypnotic induction techniques is quite understandable, since we know that any conscious experience may trigger a context (4.xx). Hypnosis involves a context, one of minimal competition for GW access. Relaxation and trust for the hypnotist may be simply different ways of describing this absorbed state.

We can also make sense of some of the negative results, the

absence of expected features of hypnosis. Hypnosis appears to be such a spectacularly different state of mind that many‹j ‹ researchers expected to find major physiological and personality differences. But if we assume that hypnosis is not an unusual state at all, but is rather a state of low competition for access to consciousness, we should find no physiological differences between hypnosis and relaxation. The absence of personality correlates is not surprising either, because we are all absorbed in our own top-most goal context, as suggested in section 7.71. In that sense all personality types involve absorption.
Finally, we should find it hard to distinguish between hypnosis and very good pretense, because successful pretense is like excellent performance in any other demanding, complex task. It requires absorption.

In sum, hypnosis may simply be ideomotor control in a state of absorption. But absorbed states are quite normal, and in a general sense, we are all absorbed in our own top-level contexts. The major difference seems to be that highly hypnotizable subjects are quite flexible in the topics of their absorption, while most people are not. Perhaps we should turn the usual question around. Instead of asking, what is îdifferentï about hypnosis? we might ask, why is flexible absorption so difficult for three-quarters of the population? What is it that is added to a "ground state" of absorption, which we all share, that resists flexible ideomotor control? We will explore this question in the next few chapters.

7.8 Conflicts between goals.

We have already discussed the possibility of competing goals and contexts (4.xx); here we will explore the implications for conflicting emotions. Goals can encounter conflict, either from other goals or from reality. All emotions involve goals, combined with real events. Happiness may result from achieving a wished-for goal; sadness involves loss of a desired object, depression is due to helplessness and hopelessness about significant life goals, anger and frustration occur when obstacles stand in the way of achieving a desired goal, fear is due to the expectation that something will happen that is fervently desired înotï to happen, love involves the goal of being with someone, and so on. All these goals can be represented in GW theory. But all these emotions involve clear, dominant goals that can be consciously achieved, delayed, thwarted, and the like.

The really difficult cases for voluntary control arise when this is not true; when there is competition for access to consciousness between different goals, so that no single goal can dominate. We have already discussed indecision due to conflicts between goals, and the possibility of an "inner argument," in which the final word wins ideomotor control. William James' discussion of "weakness of the will" and "explosive will" is also‹j ‹ relevant here, and can be treated in terms of different patterns of competing goals (1890, Chapter xx). Perhaps most intriguing, the discussion so far leads quite naturally to a viewpoint on îunconsciouslyï conflicting goals, those that may compete with the dominant goal hierarchy by generating a momentary global message that will be carried out by well-prepared systems, but with minimal metacognitive recall (7.xx above). We turn now to such unreportable goal conflicts.

7.81 A modern version of psychodynamics: Modeling unconscious goal conflict.

Unconscious conflict has been the key assumption in the long tradition of psychodynamic thought, starting with Freud and Janet in the 19th century, and continuing in an uninterrupted creative stream to the present time (Ellenberger, 1970). While it has been difficult to find solid evidence outside the clinic for many psychodynamic ideas, there is now a growing conviction among many scientific psychologists that these ideas can be tested and modeled in a reasonable cognitive framework (Baars, 1985; Erdelyi, 1985; Meichenbaum and Bowers, 19xx). This discussion is in that spirit.

7.82 Disavowed goals can be assessed by contradictions between voluntary (edited) and involuntary (unedited) expressions of the same goal.

Suppose one is furious with a friend, but finds it impossible to express this feeling. The goal hierarchy may exclude the goal of expressing anger so completely that the anger ©©© presumably some context competing for access to consciousness ©©© can only create a fleeting global goal image. Thus there will be little if any metacognitive access to the goal image. Suppose the friend asks the angry person whether he would like to meet for lunch next week, and receives the reassuring reply that "I'd like to beat you very madly" instead of "I'd like to meet you very badly." This is one kind of Freudian slip (Freud, 1901), and we have experimental evidence that deep goal conflicts can sometimes produce this kind of meaningful slip (see below). The key notion here is that we can observe an îinïvoluntary slip that expresses an emotion, but subjects will voluntarily disavow the emotion when asked about it. This may be true in general: when there is a deep conflict between goals, and one goal system dominates voluntary action and speech, it may still be possible for the excluded goal to express itself counter©voluntarily when a fleeting global goal triggers a prepared action. Voluntary‹j ‹ actions --- those that are metacognitively reportable as voluntary --- presumably have rather long-lasting goal images. Since long-lasting goal images are edited by multiple criteria, a voluntary expression of anger may be vetoed by some part of the goal hierarchy, but a fleeting angry image might gain expression if the appropriate motor systems were ready to express it. It should be edited out, but it may not be, due to a lack of editing time. All this suggests that we can use an observed contradiction between voluntary and involuntary expression of the same feeling as a signal that there is a basic goal conflict.

In general, we can suggest that emotional conflict of this kind is marked by a contradiction between voluntary and involuntary expressions of the emotion (Baars, 1985). The person makes an angry slip, but quite honestly disavows any conscious goal of expressing anger because metacognitive access to the momentary angry goal image is lost. This pattern of self™contradiction between voluntary and involuntary expressions of conflicted emotion has indeed been found with sexual slips made by males who score high on a measure of Sexual Guilt (Motley, Camden & Baars, 1979), and for angry slips in subjects who have been given a post©hypnotic suggestion of anger (Baars, Cohen & Bower, 1986). Presumably the same sort of explanation applies to the finding that female High Sex Guilt subjects show more physiological sexual arousal to an erotic tape-recording than do Low Sex Guilt subjects, even though their verbal reports show the opposite tendency (Morokoff, 1981). And Weinberger (19xx) has identified a group of "repressors" who are marked by high autonomic reactivity to emotional stimuli which they claim have no emotional effect. All these cases are marked by involuntary expression of affect along with voluntary disavowal.

Presumably conflicted subjects, such as the males who score high on Sex Guilt are in conflict between approaching and avoiding sexually desirable people (îGï and ~îGï). This conflict can be modeled as competition for access to a global workspace between goal images for avoiding and goal images for approaching sexual goals. Goal images for avoidance may encounter little competition, so that they are available longer, and are therefore reportable by relatively slow linguistic processors. But goal images for approach encounter competition from the avoidance goals, and are thus limited to very brief access to the global workspace. However, even brief access may be long enough to trigger automatic or prepared responses expressive of the forbidden goal image. The slip task presumably provides the kind of highly prepared response that allows expression to the fleeting desire to approach the attractive person.

The more these two intentions compete, the more the subject loses control over the unintentional expression of the prohibited goal, because the fleeting goal image cannot be modified as long as it is available for only a short time (7.xx). Thus the very effort to avoid thinking of the sexually attractive person may paradoxically triggers the taboo thoughts. This way of thinking‹j ‹ allows us to explain a number of phenomena that have a psychodynamic flavor, in the sense that they involve competition between contrary intentions.

These are not quite the ideas proposed by Freud, because we make no claim that deep underlying conflicts cause these phenomena --- rather, they may result from the normal functioning of the system that controls voluntary action by means of conscious goals. However, we cannot exclude the stronger Freudian hypothesis that enduring unresolved goal conflicts may initiate and bias this series of events. Indeed, the notion of a momentary conscious goal for avoidance is very similar to some versions of Freud's concept of "signal anxiety," which is sometimes said to involve a momentary experience of anxiety that signals there is something to be avoided, but without knowing what and why.

Notice that in this framework the difference between "repression" and "suppression" is only a matter of degree. If the goal image for inappropriate anger is available long enough, it may be suppressed by competition, but there will be metacognitive access to the taboo thought. But with more automaticity, or greater effort to compete against the taboo goal image, metacognitive access may be lost and we may disavow the thought quite sincerely because it is no longer accessible. However, it may still influence well©practiced action systems. Thus repression could simply be automatic suppression.

7.83 The conflict-free sphere of conscious access and control.

If goals can conflict, it makes sense to suppose that our normal, successful actions occur mostly in a domain of minimal competition. Otherwise we would show great hesitation and indecision even with acceptable and highly-practiced actions. There must be thousands of actions that are well within our power which we simply do not carry out because they conflict with other goals. Physically we are quite able to slap a close friend in the face, drop a baby, break a store window, or insult a colleague. We could deliberately break a leg, or avoid eating for a month. These possibilities rarely become conscious; they are usually not even considered. Most goals that are consciously considered are not the objects of heavy competition from other goals. Ego psychologists like Hartmann (1958) refer to this domain of minimal competition as the "conflict-free sphere of the ego."

We can easily show this conflict-free domain in our GW diagram. Figure x.xx shows an area in which the deeper and more powerful goal systems are in substantial agreement; within this domain one can get competition between local goals --- how best to get to the grocery store --- but rarely between deep goals like survival and social acceptance. (ºFootnote 3) Actions in‹j ‹ this domain are not 100% free of conflict, but they are free relative to the deep goal hierarchy. In the conflict-free domain, actions are easy to carry out, and the goal hierarchy can effectively throw its weight behind one goal or another, so that decisions can be made and adhered to with consistency (section x.xx above). The great majority of "voluntary" actions would seem to emerge from this conflict-free domain. Indeed, freedom from internal and external conflict may be a defining feature of voluntary control. It may be the psychological ground of the persistent experience of free will.

In Chapter 9 we will explore the relations between conflict- free voluntary action, and or self-attributed action. The conflict-free domain will appear as one aspect of the notion of self (9.xx).

7.9 Summary.

We began this chapter with a contrastive analysis comparing similar voluntary and involuntary actions. Next, the ideomotor theory of William James was explored and translated into GW theory; this in turn was found to explain the voluntary™involuntary contrasts. Voluntary control is treated as the result of conscious goal images that are carried out consistently with the dominant goal context; conflicts with the goal context tend to become conscious and are edited by multiple unconscious criteria. Conscious goal images are impulsive. They tend to be carried out barring competing goal images or intentions. This perspective control has implications for numerous phenomena, including slips, automaticity, psychopathological loss of control, decision making, the question of conscious access to abstract concepts, the issue of fleeting conscious events and source attribution, absorption, hypnosis, and even psychodynamics.

7.91 Some testable predictions from Model 5.

We have made some strong claims in this chapter. Not all of these are supported by direct and persuasive experimental evidence. The ideomotor theory especially needs much more support.

One approach to testing the ideomotor theory may be to use experimentally elicited slips as actions to be triggered by ideomotor goal images. A slip such as îdarn boreï - îbarn doorï may increase in frequency if one shows a rapid picture of a farm immediately before the slip. We know this is true for relatively long exposures of words related to the slip (Motley & Baars,‹j ‹

1979 a), but it may occur even if the exposure is so fast that it cannot be reported accurately, much like the Sperling figure (1.xx). A further refinement might be to evoke a conscious image immediately before the action. In that case, one might be able to study the effects of automatization of the image (x.xx). Highly automatized actions such as those studied by Shiffrin & Schneider (1977) should execute even with fleeting goal images. Finally, one might induce a cue-dependent mental image by means of post-hypnotic suggestion, with amnesia for the suggestion. Thus a highly hypnotizable subject may be told to feel an itch on his or her forehead when the experimenter clears his throat; one would expect the subject to scratch the hallucinatory itch, even though the subject was not told to scratch, merely to itch. But of course, this should generalize beyond itching and scratching. If the subject is sitting, he may be told to imagine on cue how the room looks from the viewpoint of someone who is standing up. If the ideomotor theory is correct, the subject should tend to stand up spontaneously. But since there has been no suggestion to stand up, this tendency cannot be attributed to hypnosis directly. Indeed, the tendency to stand up might be inhibited, so that one could only observe small movements in that direction, perhaps with postural muscle electrodes.

Similarly, one could induce competition against certain goal images, and study the ways in which inhibition of action can be lifted. In social situations there is a set of prohibitions against inappropriate actions, which may be induced using experimentally evoked slip techniques. If we evoked an aggressive slip directed to the experimenter, such as îyam dooï - îdamn youï, and created a distraction immediately after onset of the slip, would inhibitory restraints be lifted? If subjects were given a post-hypnotic suggestion to feel an itch on cue, but to be embarassed to scratch the itch, would the inhibition be lifted by distraction? All these techniques are potentially informative about the ideomotor hypothesis.

7.92 Some questions Model 5 does not answer.

This chapter has only addressed the issue of voluntary control: we do not yet know how this control is manifested in attention, the control of access to consciousness (8.00). Further, we do not yet have an explicit role for metacognition, which is of course necessary in order to report that some event is voluntary. And finally, we do not know why voluntary actions are always attributed to the self as agent, and why involuntary actions are not attributed to oneself (9.00). These are important questions for the following chapters.

Footnotes

1. The key role of surprise in the operational definition of voluntary control supports our previous point that expectations and intentions are quite similar. Surprise is a violation of expectations, and its existence may be used to infer the presence of an expectation. Surprise about one's own performance is similarly evidence for intentions (goal contexts), even if these intentions were previously not reportable as such. These points make sense if intentions are simply expectations about one's own performance. In GW vocabulary, they are just different kinds of contexts.

2. Dell (1986) points out that error correction can indeed occur without mismatch editing, simply by strengthening the probability of the correct plan. However, there is independent evidence from GSR studies for mismatch editing (see Baars, 1985, and in press d).

3. Of course one can design circumstances where these goal systems will be in conflict, as any dramatist knows. For

simplicity, we are here assuming that the world is stable and does not contradict the conflict-free domain.