Hall, L., Johansson, P., & de Léon, D. (2013). Recomposing the Will: Distributed motivation and computer mediated extrospection. In T. Vierkant, A. Clark & J. Kiverstein (Eds.) (2013). *Decomposing the will*. Oxford University Press: Philosophy of Mind Series. pp. 298-324.

Abstract: In this chapter we trace the problem of self-control back to its roots in research on agency and intentionality, and discuss the relationship between self-knowledge and self-control in the context of our own research on *Choice Blindness*. In addition, we provide a range of suggestions for how modern sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities, an avenue that has remained largely unexplored. In our discussion, two core concepts are introduced. The first is the concept of *Computer-Mediated Extrospection*, which builds and expands on the familiar idea of self-observation or self-monitoring as a way to gain self-knowledge. The second is the notion of *Distributed Motivation*, which follows as a natural extension of the use of precommitment and selfbinding as tools to overcome a predicted weakness of one's will.

Keywords: Choice Blindness, Intentionality, Willpower, Self-Control, Introspection, Self-Knowledge, Confabulation,

Chapter 16

Recomposing the Will: Distributed Motivation and Computer Mediated Extrospection

Lars Hall, Petter Johansson & David de Léon

ABSTRACT

In this chapter we trace the problem of self-control back to its roots in research on agency and intentionality, and discuss the relationship between self-knowledge and self-control in the context of our own research on *Choice Blindness*. In addition, we provide a range of suggestions for how modern sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities, an avenue that has remained largely unexplored.

In our discussion, two core concepts are introduced. The first is the concept of *Computer-Mediated Extrospection*, which builds and expands on the familiar idea of self-observation or self-monitoring as a way to gain self-knowledge. The second is the notion of *Distributed Motivation*, which follows as a natural extension of the use of precommitment and self-binding as tools to overcome a predicted weakness of one's will.

At the beginning of the novel *Do Androids Dream of Electrical Sheep?* by Philip K. Dick, we find Rick Deckard and his wife, Iran, in bed arguing over how to dial their daily mental states on their bedside Penfield mood organs. Deckard has wisely programmed the organ the night before to awake him in a state of general well-being and industriousness. Now he is ready to dial for the businesslike professional attitude that his electronic schedule says is needed of him today. Iran, on the other hand, has awoken to her natural proclivities and just feels irritated about Deckard's attempts to persuade her into dialing for a more productive mood. In fact, for today she has scheduled a full three-hour episode of self-accusatory depression. Deckard is unable to comprehend why anyone would ever want to willfully schedule for an episode of depression. Depression would only serve to increase the risk of her *not* using the organ at a later stage to dial into a constructive and positive mood. Iran, however, has reflected further on this dilemma and has programmed the Penfield for an automatic resetting after three hours. She will face the rest of the day in a state of "hope and awareness of the manifold possibilities open to her in the future."

In this short episode of imaginative science fiction it is not difficult to find examples of many of the most difficult conundrums of human motivation and self-control. In no small part is this of course due to Philip K. Dick being a very astute observer of the human condition, but doubtlessly it also reveals the pervasive nature of these problems in everyday life. Not being equipped with near-magical instruments of brain stimulation, people adopt all manner of strategies available to handle the ever so complicated, and in many ways both unnatural and conflicting, motivational demands of modern society. Like Deckard and Iran, how do we manage to get ourselves into the "businesslike professional attitude" that is required of us, if all we really want to do is stay in bed? Or, to up the ante, what effective, long-term means do we have to follow through on venerable goals like dieting or quitting smoking, or on general desires like becoming a more creative and lovable person? One class of answers to these questions rings particularly empty; those are the ones that in one way or another simply say, "*just do it*"—by acts of will, by showing character, by sheer motivational force, and so forth. These answers are not empty because it is difficult to find examples of people who suddenly and dramatically alter their most ingrained habits, values, and manners, seemingly without any other aid than a determined mind. It is, rather, that invoking something like "will" or "character" to explain these rare feats of mental control does little more than label them as successes. The interesting question is, rather, what we ordinary folks do when we decide to set out to pursue some lofty goal—to start exercising on a regular basis, to finally write that film script, to become a less impulsive and irritable person—if we cannot just look inside our minds, exercise our "will," and simply be done with it. The answer, we believe, is that people cope as best they can with a heterogeneous collection of culturally evolved and personally discovered strategies, skills, tools, tricks, and props. We write authoritative lists and schedules, we rely on push and pull from social companions and family members, we rehearse and mull and exhort ourselves with linguistic mantras or potent images of success, and we even set up ceremonial pseudo-contracts (trying in vain to be our own effective enforcing agencies). Often we put salient markers and tracks in the environment to remind us of, and hopefully guide us onto, some chosen path, or create elaborate scenes with manifest ambience designed to evoke the right mood or attitude (like listening to sound tracks of old *Rocky* movies before jogging around the block). We also frequently latch onto role models, seek out formal support groups, try to lock ourselves into wider institutional arrangements (such as joining a very expensive tennis club with all its affiliated activities), or even hire personal pep coaches. In short, we prod, nudge, and twiddle with our fickle minds, and in general try to *distribute* our motivation onto stable social and artifactual structures in the world.

In this chapter we trace the self-control dilemma back to its roots in research on agency and intentionality, and summarize the evidence we have accumulated in our choice-blindness paradigm for a vision of the mind as radically opaque to the self. In addition, we provide a range of suggestions for how modern sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities, an avenue that, lamentably, has remained largely unexplored. To this end, we introduce two core concepts that we hope may serve an important role in elucidating the problem of self-control from a modern computing perspective. First, we introduce the concept of *computer-mediated extrospection*, which builds and expands on the familiar idea of self-observation or self-monitoring. Second, we present the idea of *distributed motivation*, as a natural extension of previous discussions of precommitment and self-binding in the self-control literature.

Letting the Intentions Out of the Box

For someone who has a few minutes to spare for scrutinizing cognitive science–oriented flow-chart models of goal-directed behavior in humans, it would not take long to discover that in the uppermost region of the chart, a big box sits perched overlooking the flow of action. If the model deals with language, it often goes by the name of *the conceptualizer* (Levelt, Roelofts, & Meyer, 1999; Postma, 2000); if the model deals with action selection in general, it is the box containing *the prior intentions* (Brown & Pluck, 2000, but see also Koechlin & Summerfield, 2007). The reason that such an all-powerful, all-important homunculus is left so tightly boxed up in these models might simply be a reflection of our scant knowledge of how “central cognition” works (e.g., Fodor, 2000), and that the box just serves as a placeholder for better theories to come. Another more likely possibility is that the researchers often think that intentions (for action) and meaning (for language) in some very concrete sense are in the head, and that they constitute basic building blocks for any serious theory of human behavior. The line of inference is that, just because the tools of folk psychology (the beliefs, desires, intentions, decisions, etc.) are so useful, there must be corresponding processes in the brain that closely resemble these tools. In some sense this must of course be true, but the question remains whether intentions are to be primarily regarded as emanating from deep *within* the brain, or best thought of as interactive properties *of* the whole mind. The first option corresponds to what Fodor and Lepore (1993) call *intentional realism*, and it is within this framework that one finds the license to leave the prior intentions (or the conceptualizer) intact in its big, comfortable box, and in control of all the important happenings in the system. The second option sees intentional states as patterns in the behavior of the whole organism, emerging over time, and in interaction with the environment (Dennett, 1987, 1991a). Within this perspective, the question of how our intentional competence is realized in the brain is not settled by an appeal to the familiar "shape" of folk-psychological explanations. As Dennett (1987) writes:

> We would be unwise to model our serious, academic psychology too closely on these putative *illata* [concrete entities] of folk theory. We postulate all these apparent activities and mental processes in order to make sense of the behavior we observe—in order, in fact, to make as much sense as possible of the behavior, especially when the behavior we observe is our own.…each of us is in most regards a sort of inveterate auto-psychologist, effortlessly *inventing* intentional interpretations of our own actions in an inseparable mix of confabulation, retrospective self-justification, and (on occasion, no doubt) good theorizing. (91, emphasis in original)

Within this framework, every system that can be profitably treated as an intentional system by the ascription of beliefs, desires, and so forth, also *is* an intentional system in the fullest sense (see Westbury & Dennett, 2000; Dennett, 2009). But, importantly, a belief-desire prediction reveals very little about the underlying, internal machinery responsible for the behavior. Instead, Dennett (1991b) sees beliefs and desires as indirect "measurements" of a reality diffused in the behavioral dispositions of the brain/body (if the introspective reports of ordinary people suggest otherwise, we must separate the ideology of folk psychology from the folk-craft: what we actually do, from what we say and think we do; see Dennett, 1991c).

However, when reading current work on introspection and intentionality, it is hard to even find traces of the previously mentioned debate on the nature of propositional attitudes conducted by Dennett and other luminaries like Fodor and the Churchlands in the 1980s and early 1990s (for a notable recent exception, see Carruthers, 2009),[1] and the comprehensive collections on folk psychology and philosophy of mind from the period (e.g., Bogdan, 1991; Christensen & Turner, 1993) now only seem to serve as a dire warning about the possible fate of ambitious volumes trying to decompose the will!

What we have now is a situation where "modern" accounts of intentionality instead are based either on concepts and evidence drawn from the field of motor control (e.g., emulator/comparator models; see Wolpert & Ghahramani, 2004; Grush, 2004) or are is built almost purely on introspective and phenomenological considerations. This has resulted in a set of successful studies of simple manual actions, such as pushing buttons or pulling joysticks (e.g., Haggard, Clark, & Kalogeras, 2002; Moore, Wegner, & Haggard, 2009; Ebert & Wegner, 2010), but it remains unclear whether this framework can generalize to more complex and long-term activities. Similarly, from the fount of introspection some interesting conceptual frameworks for intentionality have been forthcoming (e.g., Pacherie, 2008; Gallagher, 2007; Pacherie & Haggard, 2010), but with the drawback of introducing a bewildering array of “senses” and “experiences” that people are supposed to enjoy. For example, without claiming an exhaustive search, Pacherie’s (2008) survey identifies the following concepts in need of an explanation: “awareness of a goal, awareness of an intention to act, awareness of initiation of action, awareness of movements, sense of activity, sense of mental effort, sense of physical effort, sense of control, experience of authorship, experience of intentionality, experience of purposiveness, experience of freedom, and experience of mental causation” (180).

While it is hard to make one-to-one mappings of these “senses” to the previous discussion of intentional realism, the framework of Dennett entails a thorough skepticism about the deliverances of introspection, and if we essentially come to know our minds by applying the intentional stance toward ourselves (i.e., finding out what we think and what we want by interpreting what we say and what we do), then it is also natural to shift the focus of agency research away from speculative senses and toward the wider external context of action. From our perspective as experimentalists, it is a pity that the remarkable philosophical groundwork done by Dennett has generated so few empirical explorations of intentionality (see Hall & Johansson, 2003, for an overview). This is especially puzzling because the counterintuitive nature of the intentions-as-patterns position has some rather obvious experimental implications regarding the fallibility of introspection and possible ways to investigate the nature of confabulation. As Carruthers (2009) puts it: “The account . . . predicts that it should be possible to induce subjects to *confabulate* attributions of mental states to themselves by manipulating perceptual and behavioral cues in such a way as to provide misleading input to the self-interpretation process (just as subjects can be misled in their interpretation of others)” (123).

Choices That Misbehave

Recently, we introduced *choice blindness* as a new tool to explicitly test the predictions implied by the intentional stance (Johansson et al., 2005). Choice blindness is an experimental paradigm inspired by techniques from the domain of close-up card magic, which permits us to surreptitiously manipulate the relationship between choice and outcome that our participants experience. The participants in Johansson et al. (2005) were asked to choose which of two pairwise presented female faces they found most attractive. Immediately after, they were also asked to describe the reasons for their choice. Unknown to the participants, on certain trials, a double-card ploy was used to covertly exchange one face for the other. Thus, on these trials, the outcome of the choice became the opposite of what they intended (see figure 16.1).

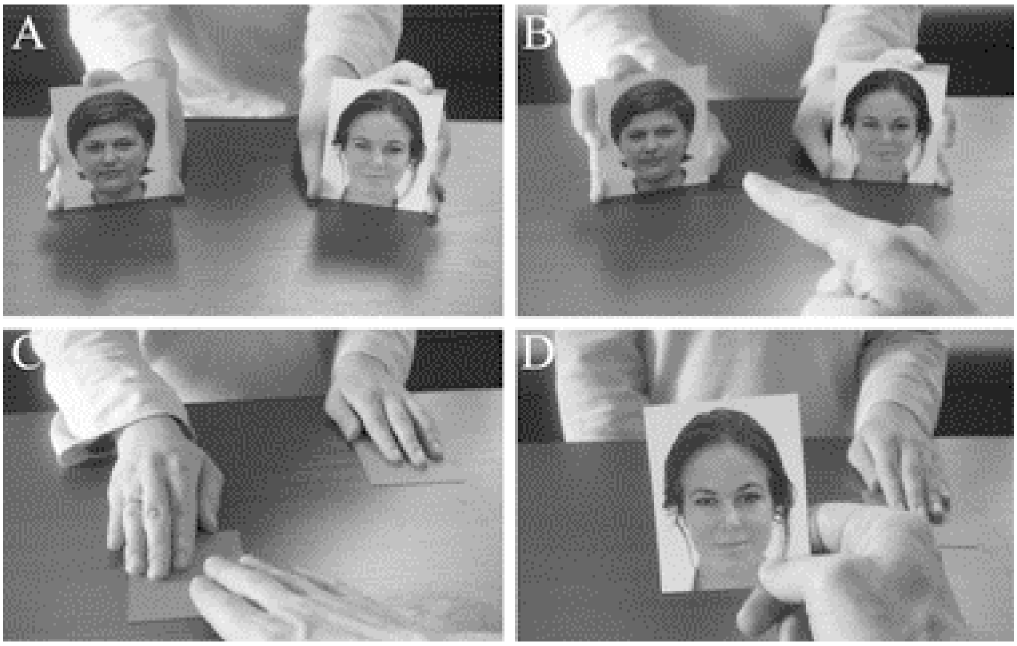

Figure 16.1 A snapshot sequence of the choice procedure during a manipulation trial. (A) Participants are shown two pictures of female faces and asked to choose which one they find most attractive. Unknown to the participants, a second card depicting the opposite face is concealed behind the visible alternatives. (B) Participants indicate their choice by pointing at the face they prefer the most. (C) The experimenter flips down the pictures and slides the hidden picture over to the participants, covering the previously shown picture with the sleeve of his moving arm. (D) Participants pick up the picture and are immediately asked to explain why they chose the way they did.

From a commonsense perspective it would seem that everyone immediately would notice such a radical change in the outcome of a choice. But that is not the case. The results showed that overall the participants detected less than 75 percent of the manipulated trials, while nevertheless being prepared to offer introspectively derived reasons for why they chose the way they did. An extensive debriefing procedure was used after the experiment to make sure that the participants who had shown no signs of detection actually were unaware of the manipulation. When we told the participants that we had in fact switched the pictures, they often showed great surprise, even disbelief at times, which indicates that they were truly unaware of the changes made during the experiment.[2]

When analyzing the reasons the participants gave, it was clear that they often confabulated their answers, as when they referred to unique features of the previously rejected face as being the reason for having made their choice (e.g., stating, “I liked the earrings” when the option they actually preferred did not have any). Additional analysis of the verbal reports in Johansson et al. (2005) as well as Johansson et al. (2006) also showed that very few differences could be found between cases where participants talked about a choice they actually made and those trials where the outcome had been reversed. One interpretation of this is that the lack of differentiation between the manipulated and nonmanipulated reports cast doubt on the origin of the nonmanipulated reports as well; confabulation could be seen to be the norm, and “truthful” reporting something that needs to be argued for.

We have replicated the original study a number of times, with different sets of faces (Johansson et al., 2006), for choices between abstract patterns (Johansson, Hall, & Sikström, 2008), and when the pictures where presented onscreen in a computer-based paradigm (Hall & Johansson, 2008). We have also extended the choice-blindness paradigm to cover more naturalistic settings, and to attribute- and monetary-based economic decisions. First, we wanted to know whether choice blindness could be found for choices involving easily identifiable semantic attributes. In this study participants made hypothetical choices between two consumer goods based on lists of general positive and negative attributes (e.g., for laptops: low price, short battery-life, etc.), and then we made extensive changes to these attributes before the participants discussed their choice. Again, the great majority of the trials remained undetected (Johansson et al., in preparation). In a similar vein, we constructed a mock-up version of a well-known online shopping site and let the participants decide which of three MP4 players they would rather buy. This time we had changed the actual price and memory storage of the chosen item when the participants reach the “checkout” stage, but despite being asked very specific questions about why they preferred this item and not the other, very few of these changes were detected (Johansson et al., in preparation). Second, we have also demonstrated the effect of choice blindness for the taste of jam and the smell of tea in an ecologically valid supermarket setting. In this study, even when participants decided between such remarkably different tastes as spicy cinnamon-apple and bitter grapefruit, or between the sweet smell of mango and the pungent Pernod, was less than half of all manipulation trials detected (Hall et al., 2010). This result shows that the effect is not just a lab-based phenomenon; we may display choice blindness for decisions made in the real world as well.

Since the publication of Johansson et al. (2005), we have been repeatedly challenged to demonstrate that choice blindness extends to domains such as moral reasoning, where decisions are of greater importance, and where deliberation and introspection are seen as crucial ingredients of the process (e.g., Moore & Haggard, 2006, commenting on Johansson et al., 2006; see also the response by Hall et al., 2006). In order to meet this challenge, we developed a magical paper survey. In this experiment, the participants were given a two-page questionnaire attached to a clipboard and were asked to rate to what extent they agreed with either a number of formulations of fundamental moral principles (such as: “Even if an action might harm the innocent, it is morally permissible to perform it,” or “What is morally permissible ought to vary between different societies and cultures”), or morally charged statements taken from the currently most hotly debated topics in Swedish news (such as: “The violence Israel used in the conflict with Hamas was morally reprehensible because of the civilian casualties suffered by the Palestinians,” or “It is morally reprehensible to purchase sexual services even in democratic societies where prostitution is legal and regulated by the government”). When the participants had answered all the questions on the two-page form, they were asked to read a few of the statements aloud and explain to the experimenter why they agreed or disagreed with them. However, the statements on the first page of the questionnaire were written on a lightly glued piece of paper, which got attached to the backside of the survey when the participants flipped to the second page. Hidden under the removed paper slip was a set of slightly altered statements. When the participants read the statements the second time to discuss their answers, the meaning was now reversed (e.g., “If an action might harm the innocent, it is morally *reprehensible* to perform it,” or “The violence Israel used in the conflict with Hamas was morally *acceptable despite* the civilian casualties suffered by the Palestinians”). Because their rating was left unchanged, their opinion in relation to the statement had now effectively been reversed. Despite concerning current and well-known issues, the detection rate only reached 50 percent for the concrete statements, and even less for the abstract moral principles

We found an intuitively plausible correlation between level of agreement with the statement and likelihood of detection (i.e., the stronger participants agreed or disagreed, the more likely they were to also detect the manipulation), but even manipulations that resulted in a full reversal of the scale sometimes remained undetected. In addition, there was no correlation between detection of manipulation and self-reported strength of general moral certainty.

But perhaps the most noteworthy finding here was that the participants that did not detect the change also often constructed detailed and coherent arguments clearly in favor of moral positions they had claimed that they did not agree with just a few minutes earlier (Hall et al., in press). Across all conditions, not counting the trials that were detected, 65 percent of the remaining trials were categorized as strong confabulation, with clear evidence that the participants now gave arguments in favor of the previously rejected position.

We believe the choice-blindness experiments reviewed here are among the strongest indicators around for an interpretative framework of self-knowledge for intentional states, as well as a dramatic example of the nontransparent nature of the human mind. In particular, we think the choice-blindness methodology represents a significant improvement to the classic and notorious studies of self-knowledge by Nisbett and Wilson (1977; see Johansson et al., 2006). While choice blindness obviously puts no end to the philosophical debate on intentionality (because empirical evidence almost never settles philosophical disputes of this magnitude; Rorty, 1993), there is one simple and powerful idea that springs from it. Carruthers (2009) accurately predicted that it would be possible to “induce subjects to confabulate attributions of mental states to themselves by manipulating perceptual and behavioral cues in such a way as to provide misleading input to the self-interpretation process” (123), but there is also a natural flip side to that prediction—if our systems for intentional ascription can be fooled, *then they can also be helped!* If self-interpretation is a fundamental component in our self-understanding, it should be possible to augment our inferential capacities by providing more and better information than we normally have at hand.

To this end, in the second section of this chapter, we introduce *computer-mediated extrospection* and *distributed motivation* as two novel concepts inspired by the Dennettian view. For intentional realists, if there is anything in the world that our private introspections tell us with certainty, it is what we *believe*, *desire*, and *intend* (Goldman, 1993). From this perspective, it would seem that a scheme of capturing and representing aspects of user context, *for the supposed benefit of the users themselves*, would be of limited value. Such information would at best be redundant and superfluous, and at worst a gross mischaracterization of the user's true state of mind. However, we contend, this is exactly what is needed to overcome the perennial problem of self-control.

THE FUTURE OF SELF-CONTROL

Computer-Mediated Extrospection

In our view, one of the most important building blocks to gain reliable knowledge about our own minds lies in realizing that it often is a mistake to confine judgment of self-knowledge to a brief temporal snapshot, when the rationality of the process instead might be found in the distribution of information traveling *between* minds: in the asking, judging, revising, and clarifying of critical, communal discourse (Mansour, 2009). As Dennett (1993) says: "Above the biological level of brute belief and simple intentional icons, human beings have constructed a level that is composed of *objects* that are socially constructed, replicated, distributed, traded, endorsed ("I'll buy that!"), rejected, ignored, obsessed about, refined, revised, attacked, advertised, discarded" (230). The point about critical communal discourse as a basis for making better self-ascriptions also naturally extends to the use of new tools and technologies to improve our self-understanding. Studies have shown that if people are simply asked to introspect (about their feelings, about the reasons for their attitudes, about the causes of their behavior, etc.), they often end up with worse judgments than the ones they initially provided (Wilson & Dunn, 2004; Silvia & Gendolla, 2001; Dijksterhuis & Aarts, 2010). On the other hand, when people are given an enhanced ability to *observe* their own behavior, they can often make sizable and profitable revisions to their prior beliefs about themselves (e.g., by way of video capture in social interaction and collaboration; see Albright & Malloy, 1999). For example, Descriptive Experience Sampling (DES) is said to be an introspective research technique. It works by using a portable beeper to cue subjects at random times, "to pay immediate attention to their ongoing experience at the moment they heard the beep. They then jot down in a notebook [or PDA] the characteristics of that particular moment" (Hurlburt & Heavey, 2001, 400; for other similar techniques, see Scollon, Kim-Prieto, & Diener, 2003; Christensen et al., 2003). Later, an in-depth interview is conducted in which the experiences are elaborated upon. What is interesting is that most participants when confronted with the processed data from the sampling protocols are surprised by some aspects of the results (e.g., Hurlburt & Heavey, 2001, describe a case of a man named Donald who discovers in the protocols that he has frequent angry thoughts directed at his children, something he was completely unaware of before). Similarly, by the use of external DES-like probes in the study of task-unrelated thought (TuT, or simply "mind wandering"), it has repeatedly been shown that participants underestimate how much their minds tend to wander—that is, that they are often *unaware* of zoning out from the task at hand (Smallwood & Schooler, 2006; Smallwood, McSpadden, & Schooler, 2008 Smallwood, Nind, & O'Connor, 2009; Christoff et al., 2009, an effect that can be tied to practical consequences outside the lab, such as educational or occupational goals (McVay, Kane, & Kwapil, 2009; Smallwood, Fishman, & Schooler, 2007; but see Baars, 2010).

Most important for us, even if the particular theories about introspection at play here are contested (e.g., see the discussion in Hurlburt & Schwitzgebel, 2007, or the exchange between Smallwood & Schooler, 2006, and McVay & Kane, 2010), there is an undeniable power for self-discovery in the external tools that enable the systematic gathering and processing of the data.[3]

But why stop with a single impoverished channel of verbal reports, when we can use modern technology to sense and compile a fantastic array of data about ourselves? The ubiquitous vision is one in which computers take an increasing part in our everyday activities, in ways that mesh naturally with how people think, act, and communicate (Bell & Dourish, 2007; Greenfield, 2006; Poslad, 2009). Work within ubiquitous computing and context awareness has made us increasingly familiar with computers that mediate our interactions with the world, *but what about computers that mediate our interactions with ourselves?* In the same manner that computers can be made more powerful by letting them gain information about the user, we also believe users can be made smarter and more powerful by letting them gain additional knowledge about themselves.

In a pioneering effort, Gordon Bell in the MyLifeBits project (see Gemmel et al., 2002; Gemmell, Bell, & Lueder, 2006; Bell & Gemmel, 2009) has collected and digitized every conceivable aspect of his own life over the span of several years. Similarly, but with an even denser assortment of wearable sensors, Clarkson (2002) gathered around-the-clock measurements over several weeks. Apart from the obvious implications for remembrance, this allows a powerful form of *personal data mining* that can reveal interesting, unintuitive, and predictive patterns in our everyday behavior. An even more ambitious approach is that of Roberts (2004, in 2010, who gathered data about himself for two decades (concerning sleep, weight loss, cognitive acuity, etc.) and subjected it to a quasi-experimental approach to overcome obstacles and improve his lot. These are three examples from a rapidly growing public trend in augmenting our inferences and attributions with extensive tracking of self-data (e.g., see the portal at http://www.quantifiedself.com/, or the services at http://your.flowingdata.com/ or http://daytum.com/, which are specifically geared the toward quantification and data mining of information gathered about the self). We believe this type of observation—what we call *computer-mediated extrospection* (CME)—is a very promising domain to explore, and that it holds great potential for improving our self-knowledge, and to extend our powers of self-regulation and control.

Drawing upon existing research in ubiquitous computing (and from conceptual neighbors like wearable computing, telemedicine, affective computing, and persuasive computing), it can be seen that capturing user context occupies center stage in human-computer interaction (Dey, Abowd, & Salber, 2001). The typical and most easily accessible context for CME is that of macrolevel activity markers, classified on a physical, intentional, and even interactive-social level (e.g., see Dalton & O'Laighin, 2009; Bajcsy et al., 2009). But perhaps even more interesting from a CME perspective are the more "intimate" measures that can be gathered from medical and/or psychophysiological monitoring. Recently, an explosion in the field of wireless, wearable (or, in some cases, even off-body) sensing has enabled reliable measuring of (among other things) electrocardiogram, blood pressure, body/skin temperature, respiration, oxygen saturation, heart rate, heart sounds, perspiration, dehydration, skin conductivity, blood glucose, electromyogram, and internal tissue bleeding (for an overview, see Pantelopoulos & Bourbakis, 2010; Kwang, 2009; Frantzidis et al., 2010). It is from these sensors, and in particular from wireless, dry electroencephalogram (EEG; Gargiulo et al., 2008; Chi & Cauwenberghs, 2010), that it is possible to build up the most critical CME variables, such as the detection and continuous monitoring of arousal, vigilance, attention, mental workload, stress, frustration, and so on (see Pan, Ren, & Lu, 2010; Ghassemi et al., 2009; Henelius et al., 2009; Grundlehner et al., 2009).

Distributed Motivation

As we stated in the opening paragraphs, the problem of self-control is not just a problem manifested in the behavior of certain “weak-willed” individuals, and it is not only operative in such salient and life-threatening domains as craving and addiction, but also in the minute workings of everyday plans, choices, and actions. Ameliorative action is as pertinent to the dreadful experience of withdrawal from heroin as it is to innocuously hitting the snooze button on the alarm clock and missing the first morning bus to school (Rachlin, 2000; Ainslie, 2001). Maglo, Gollwitzer, and Oettingen (this volume) present the evidence for the effectiveness of (so-called) implementation intentions (IMPs), which has shown that when people are prompted to elaborate a long list of very specific contingency goals (of the form “when situation X arises, I will perform response Y”), they are also significantly more likely to perform that action (Gollwitzer, 1999; Webb & Sheeran &, 2008). This effect has been repeatedly demonstrated in real-world environments, for example, in relation to rehabilitation training after surgery, to keeping up an exercise program, to eating more healthy food, to breast self-examination and screening for cervical cancer (see Gollwitzer & Sheeran, 2006, for a recent meta-analysis, but see also Sniehotta, 2009, Wood & Neal 2007. But why does forming IMPs work? Is it not enough to have “normal” intentions to act accordingly? Maglio, Gollwitzer, and Oettingen (this volume) favor the explanation that IMPs “create instant habits” and “pass the control of one’s behavior to the environment” (Gollwitzer, 1999), and they choose to frame their discussion of IMPs around the well-known parable of Odysseus and the Sirens. They write:

> In the service of [Odysseus’] goal, he consciously willed an explicit plan—having himself tied to the mast of his ship. From there, however, he had in a sense surrendered his conscious intent to nonconscious control: though his conscious will had changed (e.g., to succumb to the temptation of the Sirens),

> the bounds of the rope remained, guiding his behavior without his conscious intent. From our perspective, the rope provides a simple metaphor for the form and function of planning that specifies when, where, and how to direct action control in the service of long-term goals. (p. XXX).

Indeed, like Odysseus facing the Sirens we often know that we will find ourselves in conditions where we are likely to do something detrimental to our long-term goals, and like Odysseus tying himself to the mast we would often like to be able to *self-bind* or *precommit*, and avoid or resist such temptations. As in the episode from *Do Androids Dream of Electrical Sheep?*, when Deckard chooses to have his Penfield awake him in an industrious mood to avoid the lure of the warm bed, and Iran programs an automatic resetting to block the self-perpetuating nature of the induced depression, we would often like to be able to choose our course of action in a calm moment of reflection rather than having to battle it out in the grip of powerful urges.

For all the practical potential of IMPs, we think it is a disservice to place them next to the mighty Odysseus. The Greek king adventurer was truly and effectively bound at the mast, but Gollwitzer himself admits that IMPs "need to be based on *strong* goal intentions. As well, certain types of implementation intentions work better than others, and people need to be *committed* to their implementation intentions" (Gollwitzer, 1999, 501, our emphasis). One might reasonably wonder why we need the extra "old-school" willpower that allows us to entertain "strong" goal intentions, and be "committed" to our implementation intentions, when the whole idea of the concept was to relieve us of the burden to consciously initiate action in the face of temptations and distractions. In fact, looking at the literature, it is clear that IMPs face a disturbing creep of "moderating" variables—they are less effective for more impulsive participants (Churchill & Jessop, 2009), they only work for people with high self-efficacy (Lippke et al., 2009), they are curtailed by preexisting "response biases" (Miles &

Proctor, 2008), “habit strength” (Webb, Sheeran, & Luszczynska, 2009), as well as the “stability” of the intentions (Godin et al., 2010) and the strength of the “goal desires” (Prestwich, Perugini, & Hurling, 2008). In addition, IMPs are generally only effective when they are provided by the experimenter, who has an expert knowledge of the (often controlled) stimuli and contingencies the participants will encounter (Sniehotta, 2009). In relation to this, the obvious question is, why settle for implementation intentions as a metaphor for Odysseus and the Sirens. Why not implement the *actual* strategy of external binding?

This is what we try to capture with our second concept *distributed motivation*: the general strategy of using stable features of both the social and the artifactual environment to scaffold the process of goal attainment. As such, distributed motivation is a subclass of the well-established theory of distributed cognition (Hutchins, 1995; Clark, 2008; Hollan, Hutchins & Kirsh, 2000). Distributed cognition deals with computational processes distributed among agents, artifacts, and environments. It is a set of tools and methodologies that allow the researcher to look beyond simple “cognizant” agents and shift the unit of analysis to wider computational structures. As previewed in our discussion of Maglio, Gollwitzer, and Oettingen (this volume), one of the most central features of our notion of distributed motivation is the concept of *precommitment* or self-binding. The tale of Odysseus and the Sirens is a standard illustration of this principle (Elster, 2000; for an in-depth treatment, see Sally, 2000a, 2000b). What we would like to argue here is that the image of the clever Odysseus foiling the Sirens might serve as a promising template for the design of modern remedies based on ubiquitous and context-aware technology. While people generally strive to approximate the Odyssean ideal in their daily self-regulation behavior, they seldom manage to create conditions of precommitment stable enough to sustain them through complex and difficult problems. As sure as the fact that the majority of folk strategies of self-control have been tried and tested in harsh conditions of cultural evolution, or over the full life span of incessantly extrospecting individuals, and that they embody considerable pragmatic wisdom, is also the fact that they fail miserably when looked at on a societal scale.

The problem with most folk strategies is of course that *they do not have enough binding power* (sadly, the injunctions are often no stronger than the glue on the back of the post-it notes they are written on). For example, an often-told anecdote in the context of research on self-control is that of the young African American man that made a “powerful” commitment to pay US$20 to the Ku Klux Klan every time he smoked a cigarette. In contrast to many other cases, it is easy to understand the force this commitment *might* have on his behavior, but the fact still remains that once he has succumbed to the temptation, nothing really compels him to transfer money to the KKK. But if no such crucial deterrent for future behavior can be established, then why on earth should he adjust his behavior in relation to the commitment to begin with? Without going into philosophical niceties, it is easy to see that there is something deeply paradoxical about this kind of self-punishment. Indeed, if one really could exert the type of mental control that effectively *binds* oneself to pay the smoking fee to the KKK, then why not just simply bind oneself not to smoke in the first place?

However, even something as lowly as a pigeon can act in a self-controlled manner in a suitably arranged environment. Given a choice between pecking an illuminated button, and be administered one morsel of food after 10 seconds of delay, or pecking another button to receive twice as much after 14 seconds of delay, pigeons strongly prefer the second alternative (if the rewards were equally large, they would of course go for the one with the shorter delay). Since the pigeons clearly value the second alternative more, they should continue to do so up until the time of delivery. However, this is not always the case. With a simple manipulation of the reward contingencies it is possible to induce “irrational” choice behavior. If the pigeons are presented with the same choice pair, but given an opportunity to “reconsider” after 10 seconds (i.e., the buttons are illuminated again to allow a peck to discriminate between one unit immediately, or two units after an additional 4 seconds), the pigeons switch to the immediate and lesser reward (Rachlin, 2000). What is irrational about this? one may ask. Are pigeons not allowed to change their minds? Well, of course they are, but the poor pigeons who live in a laboratory that has the “tempting” reconsideration-button installed will award themselves considerably less food than their friends down the hall. In fact, in some sense, the pigeons seem to “realize” this. If yet another choice-button is introduced in the experiment, this time giving the pigeons a chance to eliminate the reconsideration-button (i.e., a peck on the new button prevents the reconsideration option from being illuminated), they consistently choose to do so (Rachlin, 2000). Thus, the pigeons show self-control by precommitment to their earlier choice. What is so remarkable about this example is that pigeons are manifestly not smart. Instead, it is clear that the intelligence of the system lies as much in the technology of the setup as in the mechanisms of the pigeon’s nervous system.

In the following sections we discuss how the conceptual tools we have proposed (*CME* and *distributed motivation*) can be applied and tailored to the demands of particular self-control problems. We start with comparatively less difficult problems and move on to harder ones.

CME and Distributed Motivation in Action

Self-Monitoring

The starting point for many discussions of self-control is the observation that people are often aware of their self-control problems but seldom *optimally aware* of the way these problems are expressed in their behavior, or under what contingencies or in which situations they are most prone to lapses in control (what is called *partial naïveté* in behavioral economics). Most likely, this is due to a mix of biased self-perception, cognitive limitations, and lack of inferential activity (Frederick, Loewenstein, & O'Donoghue, 2002). Within this domain, we see two rough categories of CME tools that could serve to correct faulty self-perceptions.

First, CME can capture and represent information that we normally successfully access and monitor, but which we sometimes *momentarily* fail to survey. The phenomenology of self-control lapses is often completely bereft of any feeling of us having consciously weighed alternatives and finally chosen the more tempting one. Instead, we often just find ourselves, post hoc, having completed an action that we did not previously intend to do (Elster, 2000; Ainslie, 2001). Studies have shown that while humans are quite capable at self-monitoring when given clear directives and timely external prompts, performance quickly deteriorates under natural conditions (Rachlin, 2000; Schooler, 2002; Smallwood & Schooler, 2006). (Compare not trying to scratch an itch under stern scrutiny in the doctor's office, and not scratching it later while watching TV.) The degree of self-monitoring, in turn, greatly influences the nature of our self-control behavior. There is a big difference between smoking a cigarette that happens to be the 24th of the day and being aware that one is about to light up the 24th cigarette for the day. The simple fact of providing accurate monitoring of self-control-related context has been shown to markedly reduce the incidence of self-control lapses (Rachlin, 2000; Fogg, 2003). The problem is of course that it is almost as difficult to stay constantly vigilant and attentive to such context as it is to control the behavior in the first place. This, we surmise, is an area where the use of context-aware technology and CME would be of great use (see Quinn et al. 2010, for a recent and powerful example of CME of bad habits).

Second, instead of helping people to monitor what they are doing right now, CME could be used to predict what they are just about to do. By using more intimate contextual measures like the psychophysiological state of the user, these micro-predictions should be situated at the moment of activity, and come (minutes or seconds) *before* the actual action is performed. For some types of self-control problems this will be comparatively easy. For example, any goals having to do with strong emotions (like trying to become a less aggressive person or trying to stifle unproductive anger in marital disagreements) will be an ideal target for CME micro-prediction. As Elster (2000) has pointed out, advice about emotion regulation most often fails simply because it comes *after* the unwanted emotion has already been aroused and taken full effect upon behavior. At an earlier stage such advice might have been perfectly effective (i.e., here the proper assessment of the need for self-control is as important as the control itself). Considerable research already exists on psychophysiological markers that indicate the implicit buildup or expression of emotional states not only for anger and aggression but also for more subtle conditions like frustration, stress, and anxiety (e.g., Belle et al., 2010; Hosseini & Khalilzadeh, 2010). Promising efforts have also been made to identify similarly predictive profiles for less obviously emotional behavior like smoking and gambling (Parker & Gilbert, 2008; Goudriaan et al., 2004). To increase the chances of finding predictive regularities, CME technology would add an additional layer to these techniques by allowing the measurements to be individually calibrated over time and multiple contexts (Clarkson, 2002).

Active Goal Representation

In the opening discussion we cataloged some of the many cultural strategies of self-control that people employ in their daily lives and noticed how they often fail because of the lack of crucial binding power. However, degree of binding is not the only variable that determines success or failure of any particular attempt at self-control. Sometimes the solution is actually easier than we might first think. At the most basic level of analysis an often overlooked factor is the nature of the representation of the goals we are striving for. An example from the clinical literature provides a good illustration of this. Patients who have suffered damage to the prefrontal cortex (PFC) often face dramatic impairments in their ability to engage in behaviors that depend on knowledge of a goal and the means to achieve it. They distract too easily and are said to be "stimulus bound" (Miller, 2000; see also Manuck et al., 2003). Despite this, rehabilitation studies have shown that performance on difficult tasks can be fully restored to the level of control subjects on demanding clinical tasks, by the simple use of a wireless, auditory pager system that alerts the patients at random intervals to think about their goals and what they are currently doing (Manly et al., 2002; Fish et al., 2007). In this example the pager does not function as a specific memory prosthesis, like a day planner or a PDA; it is not telling the patients *what* to do. It is a cheap, global signal that tells them to think about what it was they *really wanted to do*. Similarly, for normal people, there is reason to believe that many of our common failures to follow through on goals and plans simply stem from an inability to continuously keep our goals active in the face of a bewildering array of distracting (and, of course, often tempting) stimuli. Maintenance of behavioral goals is a full-time job even for people with perfectly intact prefrontal structures (Miller & Cohen, 2001).

Thus, the first tier in any program for alleviating problems of self-control should focus on maintaining important goals in an active state. Specific types of enhancements to prospective memory exist in countless forms: from simple post-it notes, to smartphone apps that allow users to associate items or actions to be remembered with specific geographic locations (Massimi et al. 2010; see also the impressive clinical results by Berry et al. 2009, where a wearable camera from the MyLifeBits project was used to improve the memory recall of a severely amnesic patient). More general systems, like the pager system described earlier, have been far less extensively explored. This is unfortunate, because such systems could occupy an important niche that traditional remembrance agents cannot fill. What CME systems like the wireless pager promise to do is to act like a *pacemaker for the mind*, a steady signal or beacon to orient our own thinking efforts. It would not require us to specify all our actions in advance (and then give reminders to do those things), but instead encourage us to think back and apply the knowledge of our prior goals to whatever situation we happen to find ourselves in at the time of the alert (see Tobias, 2009, for a similar perspective).

A further reason to explore such applications comes from basic learning theory. Nelson and Bouton (2002; see also Bouton, 2004; Archbold, Bouton, & Nader, 2010) have found that an asymmetry exists between initial learning in any domain and subsequent attempts at unlearning such behavior (e.g., eating or drinking habits we would like to change). With few exceptions, initial learning is far less context-dependent, while attempts at unlearning generally only work in the specific context where the training took place (e.g., in a specific environment, or in a specific state of mind, or even at a specific time; see Bouton, 2004). This means that the risk of relapse is always great unless meticulous care is taken to control for contextual variables that could be of importance. Technically, this means that learning to break a bad habit does not involve unlearning the old patterns, but rather that a new form of learning has been established that (in certain contexts) inhibits the old learning. However, Nelson and Bouton (2002) have also shown that this problem can be substantially alleviated by conditioning the retraining to a salient object that is accessible in practically any context (i.e., the object in effect works as a portable context). In the light of the previous discussion, a system like the wireless pager described by Manly et al. (2002) could, with proper preparation, work both as a beacon that is used to reengage attention to our goals and simultaneously as a signal to inhibit our bad habits.

Goal Progression

As we mentioned in the earlier discussion of CME, there is a world of difference between lighting up a cigarette that happens to be the 24th of the day, and knowingly and willingly smoking the 24th cigarette of the day. But while CME technology could provide substantial help with monitoring of goals in relation to clear-cut objectives like dieting or smoking (it is a relatively straightforward task to implement context-aware devices that could count the amount of calories or cigarettes consumed), it promises to provide an even greater impact in relation to goals that are more abstract, nebulous, or distantly long-term. For example, imagine someone who has decided to become a more amiable and caring person. How would she go about fulfilling this goal, and how would she know when she has fulfilled it? One solution that is realizable by means of context-aware technology is to operationalize the goal in such a way as to be able to get discriminating feedback on the outcome of her behavior. This is a perfect job for context-aware CME technology. What computers do best is to capture, record, store, and analyze data. With the help of ubiquitous or wearable computing devices, conditions of “goal attainment” could be specified and used as an objective comparison for the agent involved. Criteria could be set in relation to any behavior, or activity, or reaction of value that can be automatically captured (number of smiles received, time spend in charity organization service, galvanic skin responses that indicate deception and lying, reductions in stress cortisol levels, environmental contexts that suggest pleasurable social interaction, number of scheduled appointments met in time, amount of empathic thoughts captured in DES etc.). But would this really capture all there is to being an amiable person? No, obviously not, but that does not detract from the fact that any change in behavior in the direction toward such a goal would be for the better. The role of CME in such cases could be seen as a form of *scaffolding* that gets people started in the direction toward some abstract or long-term goal. When the behavioral change has gained some momentum, the scaffolding can be dropped in order for more complex (and less measurable) behaviors to flourish. Another similar, but subtly different role for computational technology in monitoring goal attainment and goal criteria is provided by Ainslie (2001). He discusses the difficult problem of trying to establish self-controlled behavior by applying and following *principles*. He argues that in the cultural sphere, and over the lifetime of an individual, a natural evolution of principles takes place, such that (with very few exceptions) principles come to evolve away from what we ideally would like them to do, to instead focus on what is clear and simple and easy to uphold. That is, people who insist on keeping their goals all "in the head" often end up with very simple and impoverished goals (because how could we otherwise *remember* them; Monterosso & Ainslie, 1999). Thus, an alcoholic who is lucky enough to recover does not recover as a "social" drinker with a controlled (and presumably) positive intake of alcohol, but as one who abstains from all forms of drinking (Ainslie, 2001; see also discussion in Rachlin, 2000). Total abstinence as a principled approach is much easier to uphold because it leaves no room for subjective interpretation (a beer together with a steak is no real drink, another drink will not hurt me because I have no more cash on me, etc.), and so it does not put the user on a slippery slope. On the other hand, as Ainslie (2001, 2005) forcefully argues, what such principles completely ignore is that this situation might often not be anywhere near what the subject would really want their lives to be like. Again, what CME can bring to this situation is the promise of using computing technology to precisely measure conditions of behavior and criteria for goal attainment, in order to effectively emulate the function of principles but without having to settle for the few cases that are so clear-cut that our ordinary senses can reliably tell them apart (i.e., we could imagine that with finely tuned sensor and computing equipment, the "social" drinker could live by a CME-augmented principle that said that she is allowed to drink only once every other month, or only a certain amount each week, or only if she is at a party of a certain size, etc.).

Micro-Precommitment

While active goal representation, accurate self-monitoring, and monitoring of goal progression are important CME strategies, they are clearly less applicable in cases of genuine reward conflict. In such cases, precommitment is the right strategy to apply. On the other hand, reward conflicts come in many different flavors, and often it is not the binding power as such that determines the value of any specific scheme of precommitment. Apart from nonmetaphorical binding, what technology has to offer the age-old strategy of precommitment is a much-lowered cost and a much-increased range of operation. This is good news because some species of precommitment need to be fast and easy to set up, and should come at a very low cost. For example, we have remote controls for many electrical appliances that enable us to turn them on and off at our convenience. But we have no remotes that allow us to turn appliances off in a way that, within a set limit of time, *we cannot turn them on again* (for TV and web surfing, we have things like parental or employer control devices that can block certain channels or domains, but we have not nearly enough effective equipment for *self*-binding[4]). We can of course always climb under the sofa, pull the plug and the antenna from the TV, and put them in a place we cannot easily reach (to make TV viewing relatively inaccessible), but such ad hoc maneuvers are generally too costly and cumbersome to perform in the long run. The trick is to strike a balance between inaccessibility and flexibility. That is, for many behaviors and situations we would like to be able to make quick, easy, but transient precommitments that allow us to move beyond some momentary temptation but then expire so as not to further limit our range of alternatives. We call this *micro-precommitment* (MPC). MPC finds its primary use when the temptations we are dealing with are not overwhelming but still noticeable enough to bring us to the fall. As an example, imagine a cell phone–based location-aware system (using GPS or any other modern positioning technique) where we can instantaneously "tag" different places from which we wish to be kept. The mechanism for tagging could be as simple as having the phone in the same "cell" as the object to be tagged, or having a place-map database in the phone that allows for distance-independent blocking. Let us now say we have a minor shoe-shopping compulsion and walk around town on an important errand. Walking down the street with this system, we could, with just a brief moment of forethought, tag an upcoming tempting shoe store. The tagging could have any number of consequences, like locking our wallet or credit card, or even tuning the store alarm to go off if we enter the premises (!). The point of MPC is *not* to set up consequences that represent maximally strong deterrents. Quite the opposite: it is a technique suited for temporarily bringing us past small but nagging distractions. Tomorrow, when we have no important errands anymore, we might want to shop for shoes again and would not want to spend our time unwinding a too forceful and elaborate precommitment scheme. In fact, since MPCs, in our view, should be as easy and cheap as possible to instigate, they *should also not be allowed to have costly or long-term consequences*.

Precommitment

If MPCs are swift and cheap and play with low stakes and short-term consequences, regular precommitment holds no such limits. For precommitment the amount of binding power and the cost of engagement are determined in relation to the magnitude of the problem and may be as strong as any agent desires. In contrast to MPC, regular precommitment should *not* come easy. To make sure that the binding represents a "true" preference, a certain amount of inertia ought to be built into any precommitment decision procedure (for a sensitive discussion of how to handle this problem, see Elster, 2000). For example, some larger casinos give patrons prone to too much gambling the option of having themselves banned from playing. Since casinos are generally equipped with rigorous security and surveillance systems, the ban can be very effectively enforced. However, one cannot just walk up to the entrance cashier and ask to be banned. The decision must be made in dialogue and with counselfrom the casino management, because once you are banned *the casino will not be coaxed into letting you in again*. As would be expected from a compulsive gambler, you soon find yourself back at the gates trying to undo your former decision. It is at this point that the casino enforces the bind by bluntly disregarding your pleas (and if the commitment was made in too light a manner, this would be an unfortunate outcome).

Craving and addiction are extremely difficult topics to approach. Behavioral abnormalities associated with addiction are exceptionally long-lived, and currently no reliable remedies exist for the pathological changes in brain-reward systems that are associated with prolonged substance abuse (Nestler, 2001; Everitt, Dickinson, & Robbins, 2001; Robinson & Berridge, 2003). With reference to precommitment, it is sometimes said that it is an ineffective strategy for handling things like addiction, because in the addicted state we supposedly never find a clear *preference platform* from which to initiate the precommitment (i.e., we do not know which of our preferences are the "true" ones). Rachlin (2000) writes: "Instead of clearly defined points of time where one strong preference gives way to its opposite we generally experience a continuous opposition of forces and apparently random alternation between making and breaking our resolutions" (54). This state of *complex ambivalence* also makes it likely that a fierce arms race will be put in motion by the introduction of any scheme of precommitment, where the addicted subject will waste precious resources and energy trying to slip through the bind of the commitment. The drug Antabuse illustrates these problems. If you take Antabuse and then have a drink, you will experience severe pain. Thus, taking Antabuse is a form of precommitment not to drink alcohol. However, alcoholics have been known to subvert the effects of the drug by sipping the alcohol excruciatingly slowly, and some even drink the alcohol despite the severe pain (Rachlin, 2000). Also, the outcome of Antabuse treatment has been generally less than satisfying because many alcoholics decide against taking the drug in the first place. In our view, this example should be taken as a cautionary tale for any overly optimistic outlook on the prospects of precommitment technology to handle really tough cases like addiction, but we do not believe it warrants a general doubt about the approach. As is evident by the fantastically prosperous industry for the supply of services and products that purport to alleviate problems of self-control (in practically any domain of life), people are willing to take on substantial commitments, in terms of time, energy, and resources, to change their current ways of life.

Take smoking as an example. What would a ubiquitous precommitment scheme for helping smokers to quit look like? First, as a foundation, some means of detecting the presence or absence of smoking-related context is needed. The context could be built from observation of the actual smoking, from traces of smoking (from smoking-related behavior patterns or from psychophysiological concomitants of smoking), and many types of sensors could be used to generate the match. For example, one sensor platform that might be used in the near future to provide robust and efficient measurement is in-blood substance detection. In relation to diabetes treatment, Tamada, Lesho, and Tierney (2002) describe a host of emerging *transdermal* (through the skin) techniques for measuring glucose levels in the blood. While not yet perfected, such sensors can be worn continually and unobtrusively by diabetics to efficiently monitor and manage their blood sugar levels. (e.g., see Gough et al., 2010). A similar system could easily be envisaged for nicotine. Yet, as many current context-aware applications have shown, a combination of many cheap and overlapping environmental sensors (i.e., things like temperature, acceleration, light, movement) might provide equally robust context measurement as a specialized subcutaneous device (Bulling, Roggen, & Troester, 2011). The great boon of ubiquitous precommitment technology is that once the basic sensing of context is in place, a multitude of distributed motivational strategies can be latched onto it, and varieties of binding can be added or subtracted depending on the nature and severity of the case. To take a dramatic example, for providing strong and relentless binding, a wireless bracelet for nicotine monitoring could be hooked up directly to the bank account of the participating subject and simply withdraw money in proportion to the amount of smoking the subject does. But to prevent loss of money, an anticipatory CME backup system that detects “lapse-critical” behavior could be employed alongside the nicotine bracelet and make automatic support calls to other participants in the program if the subject is in danger of taking a smoke. While exceptionally strong single precommitment criteria can be put in place (i.e., you lose all your money if you smoke one single cigarette), it is the possibility of mixing and merging many less forceful strategies in one system that will provide the greatest benefits. Most likely, venerable cultural strategies like situation avoidance (e.g., the shoe store “tagging” example), social facilitation, reward substitution, and so forth, will experience a strong resurgence in the hand of ubiquitous technology for distributed motivation.

Conclusion

In this chapter we discussed how the problem of self-control can be approached from a perspective on intentionality and introspection derived from the work of Dennett, and the evidence from our own choice-blindness paradigm. We have provided a range of suggestions for how sensor and computing technology might be of use in scaffolding and augmenting our self-control abilities, and we have introduce the concepts of *computer-mediated extrospection* and *distributed motivation* that we hope may serve an important role in elucidating the problem of self-control from a modern computing perspective. Some researchers have expressed pessimism about the ability of context-aware systems to make meaningful inferences about important human social and emotional states, and believe that context-aware applications can only supplant human initiative in the most carefully proscribed situations (Bellotti & Edwards, 2001). As evidenced by the current chapter, we think this pessimism is greatly overstated. Precommitment technologies offer people the option of temporary but forceful binding, aided by computer systems that will not be swayed or cajoled, and *it is through their very inflexibility that these systems have the potential to support individual self-realization*. As Dennett (2003) notes, in the domain of self-control, effectively constraining our options actually gives us more *freedom* than we otherwise would have had.

Acknowledgment

L.H. thanks the Swedish Research Council, and P.J. thanks the Bank of Sweden Tercentenary Foundation for financial support.

References

Ainslie, G. (2001). *Breakdown of will*. New York: Cambridge University Press.

Ainslie, G. (2005). Précis of Breakdown of Will. *Behavioral and Brain Sciences*, *28*(5), 635–673.

Albright, L., & Malloy, T. E. (1999). Self-observation of social behavior and metaperception. *Journal of Personality and Social Psychology*, *77*(4), 726–743.

Archbold, G. E. B., Bouton, M. E., & Nader, K. (2010). Evidence for the persistence of contextual fear memories following immediate extinction. *European Journal of Neuroscience*, *31*(7), 1303–1311.

Baars, B. J. (2010). Spontaneous repetitive thoughts can be adaptive: Postscript on "mind wandering." *Psychological Bulletin*, *136*(2), 208–210.

Bajcsy, R., Giani, A., Tomlin, C., Borri, A., & Di Benedetto, M. (2009). Classification of physical interactions between two subjects. In *BSN 2009. Sixth International Workshop on Wearable and implantable body sensor networks.*. (pp. 187–192).

Bell, G., & Dourish, P. (2007). Yesterday's tomorrows: Notes on ubiquitous computing's dominant vision. *Personal and Ubiquitous Computing*, *11*(2), 133–143.

Bell, G., & Gemmell, J. (2009). *Total recall: How the e-memory revolution will change everything.* New York: Dutton Adult

Belle, A., Soo-Yeon Ji, Ansari, S., Hakimzadeh, R., Ward, K., & Najarian, K. (2010). Frustration detection with electrocardiograph signal using wavelet transform. In *2010 International Conference on Biosciences* (pp. 91–94).

Bellotti, V. M., & Edwards, W. K. (2001). Intelligibility and accountability: Human considerations in context aware system. *Human–Computer Interaction, 16*, 193–212.

Berry, E., Hampshire, A., Rowe, J., Hodges, S., Kapur, N., Watson, P., & Browne, G. (2009). The neural basis of effective memory therapy in a patient with limbic encephalitis. *Journal of Neurology, Neurosurgery, and Psychiatry*, *80*(11), 1202–1205.

Bogdan, R. J. (Ed.). (1991). *Mind and common sense: Philosophical essays on commonsense psychology*. Cambridge: Cambridge University Press.

Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning and Memory, 11*(5), 485–494.

Brown, R. G., & Pluck, G. (2000). Negative symptoms: The “pathology” of motivation and goal-directed behaviour. *Trends in Neurosciences, 23*(9), 412–417.

Bulling, A., Roggen, D., & Troester, G. (2011). What’s in the eyes for context-awareness? *Pervasive Computing, IEEE,* April-June 2011, pp. 48-57

.

Carruthers, P. (2009). How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and Brain Sciences*, *32*, 121–182.

Chi, Y. M., & Cauwenberghs, G. (2010). Wireless non-contact EEG/ECG electrodes for body sensor networks. In *2010 International Conference on Body Sensor Networks* (pp. 297–301).

Christensen, S. M., & Turner D. R. (1993). *Folk psychology and the philosophy of mind*. Hillsdale, NJ: Erlbaum.

Christensen, T. C., Barrett, L, F., Bliss-Moreau, E., Lebo, K., & Kaschub, C. (2003). A practical guide to experience-sampling procedures. *Journal of Happiness Studies*, *4*, 53–78.

Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(21), 8719–8724.

Churchill, S., & Jessop, D. (2009). Spontaneous implementation intentions and impulsivity: Can impulsivity moderate the effectiveness of planning strategies? *British Journal of Health Psychology*, *13*, 529–541.

Clark, A. (2008). *Supersizing the mind*. Oxford: Oxford University Press.

Clarkson, B. (2002). Life patterns: Structure from wearable sensors. PhD diss., MIT

Dalton, A., & O'Laighin, G. (2009). Identifying activities of daily living using wireless kinematic sensors. In *Sixth International Workshop on Wearable and Implantable Body Sensor Networks* (pp. 87–91).

Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

Dennett, D. C. (1991a). *Consciousness explained*. Boston: Little, Brown.

Dennett, D. C. (1991b). Real patterns. *Journal of Philosophy*, *89*, 27–51.

Dennett, D. C. (1991c). Two contrasts: Folk craft versus folk science and belief versus opinion. In J. Greenwood (Ed.), *The future of folk psychology: Intentionality and cognitive science* (135–148). Cambridge: Cambridge University Press.

Dennett, D. C. (1993). The message is: There is no medium. *Philosophy and Phenomenological Research*, *53*, 889–931.

Dennett, D. C. (1994). Get real. *Philosophical Topics*, *22*, 505–568.

Dennett, D. C. (2003). *Freedom evolves.* London: Allen Lane.

Dennett, D. C. (2009). Intentional systems theory. In B. McLaughlin, A. Beckermann, & S. Walter (Eds.), *Oxford handbook of the philosophy of mind* (pp. 339–350).

Dey, A. K., Abowd, G. D., & Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human–Computer Interaction*, *16*, 167–176.

Dijksterhuis, A., & Aarts, H. (2010). Goals, attention, and (un)consciousness. *Annual Review of Psychology*, *61*, 467–490.

Ebert, J. P., & Wegner, D. M. (2010). Time warp: Authorship shapes the perceived timing of actions and events. *Consciousness and Cognition*, *19*(1), 481–489.

Elster, J. (2000). *Ulysses unbound.* Cambridge: Cambridge University Press.

Everitt, B. J., Dickinson, A., & Robbins, T. W. (2001). The neuropsychological basis of addictive behavior. *Brain Research Reviews*, *36*, 129–138.

Fish, J., Evans, J. J., Nimmo, M., Martin, E., Kersel, D., Bateman, A., & Wilson, B. A. (2007). Rehabilitation of executive dysfunction following brain injury: "Content-free" cueing improves everyday prospective memory performance. *Neuropsychologia*, *45*(6), 1318–1330.

Fodor, J. A. (2000). *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press.

Fodor, J., & Lepore, E. (1993). Is intentional ascription intrinsically normative? In Bo Dahlbom (Ed.), *Dennett and his critics* (pp. 70–82). Oxford: Blackwell.

Fogg, B. J. (2003). *Persuasive technology: Using computers to change what we think and do.* San Francisco: Morgan Kaufmann.

Frantzidis, C., Bratsas, C., Papadelis, C., Konstantinidis, E., Pappas, C., & Bamidis, P. (2010). Toward emotion aware computing: An integrated approach using multichannel neurophysiological recordings and affective visual stimuli. *Transactions on Information Technology in Biomedicine, IEEE, 14*(3), 589–597.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *40*(2), 351–401.

Gallagher, S. (2007). The natural philosophy of agency. *Philosophy Compass*, *2*(2), 347–357.

Gargiulo, G., Bifulco, P., Calvo, R., Cesarelli, M., Jin, C., & van Schaik, A. (2008). A mobile EEG system with dry electrodes. In *Biomedical Circuits and Systems Conference. BioCAS 2008. IEEE* (pp. 273–276).

Gemmell, J., Bell, G., & Lueder, R. (2006). MyLifeBits: A personal database for everything. *Communications of the ACM*, *49*(1), 88–95.

Gemmell, J., Bell, G., Lueder, R., Drucker, S., & Wong, C. (2002). MyLifeBits: Fulfilling the Memex vision. In *Proceedings of the Tenth ACM International Conference on Multimedia* (pp. 235–238).

Ghassemi, F., Moradi, M., Doust, M., & Abootalebi, V. (2009). Classification of sustained attention level based on morphological features of EEG's independent components. In *ICME International Conference on Complex Medical Engineering, 2009. CME.* (pp. 1–6).

Godin, G., Belanger-Gravel, A., Amireault, S., Gallani, M., Vohl, M., & Perusse, L. (2010). Effect of implementation intentions to change behaviour: moderation by intention stability. *Psychological Reports*, *106*(1), 147–159.

Goldman, A. (1993). The psychology of folk psychology. *Behavioral and Brain Sciences*, *16*, 15–28.

Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist*, *54*, 493–503.

Gollwitzer, P., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, *38*, 69–119.

Goudriaan, A. E., Oosterlaan, J., de Beurs, E., & Van den Brink, W. (2004). Pathological gambling: A comprehensive review of biobehavioral findings. *Neuroscience and Biobehavioral Reviews*, *28*(2), 123–141.

Gough, D., Kumosa, L., Routh, T., Lin, J., & Lucisano, J. (2010). Function of an implanted tissue glucose sensor for more than 1 year in animals. *Science Translational Medicine*, *2*(42), 42–53.

Greenfield, A. (2006). *Everyware: The dawning age of ubiquitous computing*. Berkeley, California. New Riders.

Grundlehner, B., Brown, L., Penders, J., & Gyselinckx, B. (2009). The design and analysis of a real-time, continuous arousal monitor. In *wearable and implantable body sensor networks. BSN 2009* (pp. 156–161).

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, *27*, 377–442.

Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*(4), 382–385.

Hall, L., & Johansson, P. (2003). *Introspection and extrospection: Some notes on the contextual nature of self-knowledge*. Lund University Cognitive Studies, 107. Lund: LUCS.

Hall, L., & Johansson, P. (2008). Using choice blindness to study decision making and introspection. In P. Gärdenfors & A. Wallin (Eds.), *Cognition—A Smorgasbord* (pp. 267–283). Lund, Nya Doxa.

Hall, L., Johansson, P., Sikström, S., Tärning, B. & Lind, A. (2006). How something can be said about Telling More Than We Can Know: Reply to Moore and Haggard. *Consciousness and Cognition*, *15*, 697–699.

Hall, L., Johansson, P., Tärning, B., Sikström, S., & Deutgen, T. (2010). Magic at the marketplace: Choice blindness for the taste of jam and the smell of tea. *Cognition*, *117*, 54–61.

Henelius, A., Hirvonen, K., Holm, A., Korpela, J., & Muller, K. (2009). Mental workload classification using heart rate metrics. In *Annual International Conference of the Engineering in Medicine and Biology Society* (pp. 1836–1839).

Hollan, J., Hutchins, E., & Kirsh, D. (2000) Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, *7*(2), 174–196.

Hosseini, S., & Khalilzadeh, M. (2010). Emotional stress recognition system using EEG and psychophysiological signals: Using new labelling process of EEG signals in emotional stress state. In *International Conference on Biomedical Engineering and Computer Science* (pp. 1–6).

Hurlburt, R. T., & Heavey, C. L. (2001). Telling what we know: Describing inner experience. *Trends in Cognitive Science*, *5*(9), 400–403.

Hurlburt, R. T., & Christopher L. H. (2006). *Exploring inner experience.* Amsterdam: John Benjamins.

Hurlburt, R. T., & Schwitzgebel, E. (2007). *Describing inner experience? Proponent meets skeptic.* Cambridge, MA: MIT Press.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Johansson, P., Hall, L., Kusev, P., Aldrovandi, S., Yamaguchi, Y., Watanabe, K. (in preparation). *Choice blindness in multi attribute decision making.*

Johansson, P., Hall, L., & Sikström, S. (2008). From change blindness to choice blindness. *Psychologia*, *51*, 142–155.

Johansson, P., Hall, L., Sikström, S., & Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, *310*, 116–119.

Johansson, P., Hall, L., Sikström, S., Tärning, B., & Lind, A. (2006). How something can be said about Telling More Than We Can Know. *Consciousness and Cognition*, *15*, 673–692.

Koechlin, E & Summerfield, C (2007) An information theoretical approach to prefrontal executive function. *Trends in Cognitive Science*, 11(6):229-35.

Kwang, P. (2009). Nonintrusive measurement of biological signals for ubiquitous healthcare. In *Annual International Conference of the IEngineering in Medicine and Biology Society* (pp. 6573–6575).

Levelt, W. J. M., Roelofts, A., & Meyer, A. S. (1999). A theory of lexical access in speech production. *Behavioral and Brain Sciences*, *22*(1), 1–76.

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The metacognitive error of overestimating change-detection ability. *Visual Cognition*, *7*, 397–412.

Lippke, S., Wiedemann, A., Ziegelmann, J., Reuter, T., & Schwarzer, R. (2009). Self-efficacy moderates the mediation of intentions into behavior via plans. *American Journal of Health Behavior*, *33*(5), 521–529.

Manly, T., Hawkins, K., Evans, J., Woldt, K., & Robertson, I. H. (2002). Rehabilitation of executive function: Facilitation of effective goal management on complex tasks using periodic auditory alerts. *Neuropsychologia*, *40*(3), 271–281.

Mansour, O. (2009). Group intelligence: A distributed cognition perspective. In INCOS '09 Proceedings of the International Conference on Intelligent Networking and Collaborative Systems (pp. 247–250). IEEE Computer Society Washington, DC, USA.

Manuck, S. B., Flory, J. D., Muldoon, M. F., & Ferrell, R. E. (2003). A neurobiology of intertemporal choice. In G. Loewenstein, D. Read, & R. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (139–172). New York: Russell Sage Foundation.

Massimi, M., Truong, K., Dearman, D., & Hayes, G. (2010). Understanding recording technologies in everyday life. *Pervasive Computing, IEEE*, *9*(3), 64–71.

McVay, J. C., & Kane, M. J. (2010). Does mind wandering reflect executive function or executive failure? Comment on Smallwood and Schooler (2006) and Watkins (2008). *Psychological Bulletin*, *136*(2), 188–197; discussion 198–207.

McVay, J. C., Kane, M. J., & Kwapil, T. R. (2009). Tracking the train of thought from the laboratory into everyday life: An experience-sampling study of mind wandering across controlled and ecological contexts. *Psychonomic Bulletin and Review*, *16*(5), 857–863.

Miles, J. D., & Proctor, R. W. (2008). Improving performance through implementation intentions: Are preexisting response biases replaced? *Psychonomic Bulletin and Review*, *15*(6), 1105–1110

Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience*, *1*, 59–65.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.

Monterosso, J., & Ainslie, G. (1999). Beyond discounting: Possible experimental models of impulse control. *Psychopharmacology, 146*, 339–347.

Moore, J., & Haggard, P. (2006). Commentary on "How something can be said about telling more than we can know: On choice blindness and introspection." *Consciousness and Cognition*, *15*(4), 693–696.

Moore, J. W., Wegner, D. M., & Haggard, P. (2009). Modulating the sense of agency with external cues. *Consciousness and Cognition*, *18*(4), 1056–1064.

Nelson, J. B., & Bouton, M. E. (2002). Extinction, inhibition, and emotional intelligence. In L. F. Barrett & P. Salovey (Eds.), *The wisdom in feeling: Psychological processes in emotional intelligence* (pp. 60–85). New York: The Guilford Press.

Nestler, E. J. (2001). Molecular basis of long-term plasticity underlying addiction. *Nature Reviews Neuroscience*, *2*, 119–128.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231–259.

Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, *107*, 179–217.

Pacherie, E., & Haggard, P. (2010). What are intentions? In, L. Nadel & W. Sinnott-Armstrong (Eds.), *Benjamin Libet and agency* (pp. 70–84). Oxford: Oxford University Press.

Pan, J., Ren, Q., & Lu, H. (2010). Vigilance analysis based on fractal features of EEG signals. In *International Symposium on Computer Communication Control and Automation (3CA), 2010* (pp. 446–449).

Pantelopoulos, A., & Bourbakis, N. (2010). A survey on wearable sensor-based systems for health monitoring and prognosis. *Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, *40*(1), 1–12.

Parker, A. B., & Gilbert, D. G. (2008). Brain activity during anticipation of smoking-related and emotionally positive pictures in smokers and nonsmokers: A new measure of cue reactivity. *Nicotine and Tobacco Research*, *10*(11), 1627.

Poslad, S. (2009). *Ubiquitous computing: smart devices, smart environments and smart interaction.* Wiley.

Postma, A. (2000). Detection of errors during speech production: A review of speech monitoring models. *Cognition*, *77*, 97–131.

Prestwich, A., Perugini, M., & Hurling, R. (2008). Goal desires moderate intention-behaviour relations. *British Journal of Social Psychology/British Psychological Society*, *47*, 49–71.

Quinn, J. M., Pascoe, A., Wood, W., & Neal, D. T. (2010). Can't control yourself? Monitor those bad habits. *Personality and Social Psychology Bulletin*, *36*(4), 499–511.

Rachlin, H. (2000). *The science of self-control*. Cambridge, MA: Harvard University Press.

Roberts, S. (2004). Self-experimentation as a source of new ideas: Ten examples about sleep, mood, health, and weight. *Behavioral and Brain Sciences*, *27*, 227–288.

Roberts, S. (2010). The unreasonable effectiveness of my self-experimentation. *Medical Hypotheses*. Dec;75(6):482-9.

Robinson, T. E., & Berridge, K, C. (2003). Addiction. *Annual Review of Psychology*, *54*, 25–53.

Rorty, R. (1993). Holism, intrinsicality, and the ambition of transcendence. In B. Dahlbom (Ed.), *Dennett and his critics: Demystifying mind* (pp. 184–202). Cambridge, MA: Basil Blackwell.

Sally, D. (2000a). Confronting the Sirens: Rational behavior in the face of changing preferences. *Journal of Institutional and Theoretical Economics*, *156*(4), 685–714.

Sally, D. (2000b). I, too, sail past: Odysseus and the logic of self-control. *Kyklos*, *53*(2), 173–200.

Schooler, J. W. (2002). Re-representing consciousness: Dissociations between experience and meta-consciousness. *Trends in Cognitive Science*, *6*(8), 339–344.

Scollon, C. N., Kim-Prieto, C., & Diener, E. (2003). Experience sampling: Promises and pitfalls, strengths and weaknesses. *Journal of Happiness Studies*, *4*, 5–34.

Silvia, P., & Gendolla, G. (2001). On introspection and self-perception: Does self-focused attention enable accurate self-knowledge? *Review of General Psychology, 5*(3), 241–269.

Smallwood, J., Fishman, D. J., & Schooler, J. W. (2007). Counting the cost of an absent mind: Mind wandering as an underrecognized influence on educational performance. *Psychonomic Bulletin & Review, 14*(2), 230–236.

Smallwood, J., McSpadden, M., & Schooler, J. W. (2008). When attention matters: The curious incident of the wandering mind. *Memory and Cognition*, *36*(6), 1144–1150.

Smallwood, J., Nind, L., & O'Connor, R. C. (2009). When is your head at? An exploration of the factors associated with the temporal focus of the wandering mind. *Consciousness and Cognition*, *18*(1), 118–125.

Smallwood, J., & Schooler, J. W. (2006). The restless mind. *Psychological Bulletin*, *132*(6), 946–958.

Sniehotta, F. F. (2009). Towards a theory of intentional behaviour change: Plans, planning, and self-regulation. *British Journal of Health Psychology*, *14*, 261–273.

Tamada, J. A., Lesho, M., & Tierny, M. (2002). Keeping watch on glucose. *IEEE Spectrum Online*, *39*(4), 52–57.

Tobias, R. (2009). Changing behavior by memory aids: A social psychological model of prospective memory and habit development tested with dynamic field data. *Psychological Review*, *116*(2), 408–438.

Webb, T. L., & Sheeran, P. (2008). Mechanisms of implementation intention effects: The role of goal intentions, self-efficacy, and accessibility of plan components. *British Journal of Social Psychology/British Psychological Society*, *47*, 373–395.

Webb, T. L., Sheeran, P., & Luszczynska, A. (2009). Planning to break unwanted habits: Habit strength moderates implementation intention effects on behaviour change. *British Journal of Social Psychology/British Psychological Society*, *48*, 507–523.

Westbury, C., & Dennett, D. (2000). Mining the past to construct the future: Memory and belief as forms of knowledge. In D. L. Schacter & E. Scarry (Eds.), *Memory, brain, and belief* (pp. 11–32). Cambridge, MA: Harvard University Press.

Wilson, T. D., & Dunn, E. W. (2004). Self-knowledge: Its limits, value, and potential for improvement. *Annual Review of Psychology*, *55*, 493–518.

Wolpert, D., & Ghahramani, Z. (2004). Computational motor control. In M. Gazzaniga (Ed.), *The cognitive neurosciences* (3rd ed., pp. 485–494). Cambridge, MA: MIT Press.

Wood, W., & Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological Review*, *114*(4), 843–863.

Footnotes

[1] At times, tension ran so high in this debate that one might have thought it would have been remembered for its rhetorical flair, if nothing else. As an example, Fodor and Lepore (1993) scolded Dennett's *superficialism* about the mental and professed that there really are no other ideas than commonsense "Granny-psychology" to take seriously, while Dennett (1994) in response, coined the name *hysterical realism* for Fodor's program and admitted that he regarded "the large and well-regarded literature on propositional attitudes . . . to be history's most slowly unwinding unintended *reductio ad absurdum*" (241, emphasis in original).

[2] After having probed what they thought of the experiment, and if they thought anything had felt strange with the procedure, the participants were also asked the hypothetical question if they think they *would* have noticed if we had switched the pictures. No less than 84 percent of the participants who did not detect any of the manipulations still answered that they would have noticed if they had been presented with mismatched outcomes in this way, thus displaying what might be called "choice-blindness blindness"—the false metacognitive belief of being able to detect changes to the outcome of one's choices (See Levin et al., 2000, for a similar result in relation to change blindness).

[3] Incidentally, the DES paradigm also represents one additional strong line of evidence against the concept of intentional realism. As Hurlburt (2009) writes: "As a result of 30 years of carefully questioning subjects about their momentary experiences, my sense is that trained DES subjects who wear a beeper and inspect what is directly before the footlights of consciousness at the moment of the beeps almost never directly apprehend an attitude. Inadequately trained subjects, particularly on their first sampling day, occasionally report that they are experiencing some attitude. But when

those reports are scrutinized in the usual DES way, querying carefully about any perceptual aspects, those subjects retreat from the attitude-was-directly-observed position, apparently coming to recognize that their attitude had been merely "background" or "context." That seems entirely consonant with the view that these subjects had initially inferred their own attitudes in the same way they infer the attitudes of others (150).

[4] But see the OSX self-control application by Steve Lambert (http://visitsteve.com/work/selfcontrol/), which allows the user to selectively and irrevocably (within a time limit) shut down sections of the web, or the slightly less weighty, but ever so useful Don't Dial (http://www.dontdial.com/) app for the iPhone/Android platform, which before an intoxicating evening allows the user to designate a range of sensitive phone contacts that later will be blocked from calling.