

Think, but Not Too Much: Willpower and Personal Evolution*

Carlos Alós-Ferrer[†]

This version: August 2014

Abstract

We consider a model of self-control based on dual-process theories from psychology. Decisions can be made by either an automatic/impulsive process modeling reinforcement or a controlled/deliberative one which maximizes utility. Effortful self-control (willpower) is a limited resource, i.e. an initial exertion of willpower limits the amount of willpower available for persevering later. A “personal evolution” approach shows that agents might fall into different self-control traps. One possibility is a “personal optimum” where self-control would be optimal but none is exercised. Another is excessive self-control, where willpower is wasted in initial decisions only to give in to temptation later.

(JEL D03, D83)

Keywords: Self-control, willpower, reinforcement, habit formation, dual processes.

*The author thanks Anja Achtziger, Peter Gollwitzer, Georg Kirchsteiger, Manfred Nermuth, Muriel Niederle, Fei Shi, Fritz Strack, and Alexander K. Wagner for helpful comments and discussions.

[†]Department of Economics, University of Cologne, Albertus-Magnus-Platz, D-50923 Cologne (Germany). E-mail: carlos.alos-ferrer@uni-koeln.de

1 Introduction

Self-control problems come in different flavors. Consider any task or project requiring perseverance, as dieting, sticking to an exercise plan, or writing a referee report. The typical self-control failure which comes to mind involves giving in to temptation, that is, failing to complete a task which has already been started when confronted with an alternative which appears more tempting in the short run, like eating the chocolate pie, sitting down to watch a movie, or working on your own, more interesting research. This failure can be easily conceptualized as an implementation failure of deliberative, long-run planning and a triumph of shortsighted, impulsive behavior. There are, however, self-control problems of a completely different kind. For instance, many decision makers, who have become acquainted with their own self-control capabilities over the years, argue to themselves that giving in to temptation is too likely, and hence it is better not to start at all. In this case, the problem arises from deliberation, and not in spite of it. Further, there is also some puzzling evidence that certain self-control problems might be best tackled by avoiding the conscious exercise of self-control. An extreme example can be found in the smoke-cessation literature, where it has recently been found that *unplanned* smoke-quitting attempts are twice as likely to be successful as planned ones (Ferguson et al., 2009; West and Sohal, 2006).

The fact that there are different kinds of self-control problems is well-known in applied psychology, which differentiates the shielding of an ongoing goal striving from getting started with goal striving (see e.g. Gollwitzer and Sheeran, 2006). Also well-known are some possible solutions or “strategies”. Specifically, motivation psychology has developed surprising tactics to improve self-control, whose characteristics shed light on the essence of self-control problems (see Fujita, 2011, for a recent review). The first problem, giving in to temptation, can be fought by shielding the personal goal-striving from unwanted thoughts (Achtziger, Gollwitzer and Sheeran, 2008; Bayer, Gollwitzer and Achtziger, 2010). The second problem, the failure to get started, can sometimes be counteracted by automatizing the desired behavior—in other words, thinking less about it. Compact pieces of popular wisdom also encompass this insight, from “Don’t be a quitter!” to “Just do it!” .

As an illustrative example, consider so-called “implementation intentions” (Gollwitzer, 1999). Those refer to a family of tactics arising in social psychology which can be adapted to fight different kinds of self-control problems. An implementation intention is merely an action trigger where the decision maker forms a plan in an “if/when... then...” format, specifying an anticipated cue as a condition for starting a course of action, e.g. “when I finish this proof, then I’ll start with that referee report.” Extensive evidence (see Achtziger and Gollwitzer, 2010, for a review) indicates that this produces an automatization of

behavior which actually helps attain the desired goals, by counteracting the failure-to-get-started phenomenon.¹ In a second step, implementation intentions can also be used to protect ongoing goal striving from temptation by specifying the anticipated temptations as cues, e.g. “if I am offered a snack, I’ll ask for a glass of water instead”. In the first case, the action trigger is avoiding the onset of a (possibly short-sighted) deliberative process. In the second case, it is fighting fire with fire, deliberately replacing a predictable impulsive reaction with a different one; in other words, in this case the action trigger acts as a psychological commitment mechanism, implementing the action that a previous deliberative process found optimal.²

For an economist, it might be surprising to consider that a larger involvement of deliberation versus impulsive or automatized actions might be the solution in some cases and the problem in others. Following Thaler and Shefrin (1981) and Schelling (1984), multiple selves models in economics typically address self-control problems by postulating a long-run, farsighted planner and one or several shortsighted doers which the planner attempts to control at a certain cost. For instance, in the realm of intertemporal choice (Benhabib and Bisin, 2005; Fudenberg and Levine, 2006; Ali, 2011), different utility functions are assigned to different selves, each modeling one particular motivation of a given, fixed agent. There is no question, however, that within this approach optimality would always ensue if the long-run planner could have her say. The problem is the lack of commitment mechanisms: a fully rational decision maker would like to commit to the planner’s chosen course of action.

The view from psychology is different. Dual-process theories postulate that the human mind is mainly influenced by two kinds of processes, called *automatic* and *controlled* (e.g., Bargh, 1989; Loewenstein and O’Donoghue, 2005; Evans, 2008; Weber and Johnson, 2009; Alós-Ferrer and Strack, 2014). Automatic processes capture e.g. impulsive reactions and are defined as efficient, fast, and unconscious. They rely on learned associations (brake if the traffic light is red) and often respond to simple reinforcement mechanisms. Controlled processes are the basic ingredient of deliberation and are defined as slow, partially reflected upon consciously, and consuming cognitive resources. The distinction between automatic and controlled processes is often analogous to the economists’ distinction be-

¹Implementation intentions are also helpful when the desired behavior requires a punctual action which the decision maker might delay or fail to execute under normal circumstances. For instance, in a recent field application, Milkman et al. (2011) used implementation intentions to enhance influenza vaccination rates. Nickerson and Rogers (2010) used implementation intentions to increase voter turnout in the 2008 US presidential elections.

²Other psychological manipulations follow different strategies. For instance, “cognitive reconstrual” aims for a reappraisal of the decision framework to promote a more abstract evaluation which should reduce impulsive behavior. Examples range from Moore, Mischel and Zeiss (1976) to Fujita and Han (2009).

tween a heuristic/intuitive and a rationality benchmark (Kahneman, 2003). In a second step, dual-process accounts of human behavior are often extrapolated to “dual-system” theories, which postulate that different information-processing systems are responsible for handling different types of processes, e.g. an impulsive “System I” and a reflective “System II”, as recently popularized by Kahneman (2011) (see also Strack and Deutsch, 2004). Loewenstein and O’Donoghue (2005) label the two systems “affective” and “deliberative”, while Epstein (1994) calls them “experiential” and “rational”. The dual-process approach has also been proposed as a useful formalization of intuitive vs. deliberative decision making, e.g. in managerial contexts (Dane and Pratt, 2007; Achtziger and Alós-Ferrer, 2014).³

The multiple selves approach has been inspired by dual-process theories. There are, however, fundamental differences which also distinguish the model presented here from the multiple selves literature. In multiple-selves models, every self is “rational” in the sense of being a utility-maximizer given a certain utility function. Hence, they amount to intrapersonal games. This is not the case for dual-process theories. Automatic processes are meant to capture impulsive reactions and stimulus-response phenomena. Accordingly, they might not even be rationalizable as utility-maximizing actors, whether shortsighted or not. Rather, they might capture much simpler processes as e.g. pure reinforcement. Further, controlled processes are also not assumed to correspond to an all-knowing rational planner (in spite of the occasional use of the adjective “rational” to refer to them). For instance, if a fully rational planner decides that it is better not to start a task because of anticipated temptation, quitting must certainly be optimal. This is not the case if the decision to quit is the result of deliberation as captured by a controlled process, but that deliberation is shortsighted—a controlled process will, more often than not, correspond to myopic optimization of a short-run objective function.

The main objective of this paper is to provide a simple, parsimonious model of self-control specifically based on the existence of two different systems of thinking. That is, we deliberately move away from the intrapersonal game approach in favor of ideas from the psychological dual-process literature, with the objective of accounting for different self-control failures, including both “impulsive” giving in to temptations and “deliberative” failures to get started. While we will rely on intuitions and insights shared with multiple-selves models, we depart from the planner-doer formalism in favor of an approach

³Evidence from social neuroscience (Lieberman, 2003; Lieberman, Jarcho and Satpute, 2004) suggests that processes of different types might be associated with different neurological systems in the human brain. For instance, evidence of this type was provided by McClure et al. (2004) in the realm of intertemporal decision making and by Sanfey et al. (2003) in the context of an ultimatum game. However, the view that such neural systems are clearly separated might be too simplistic (Frank, Cohen and Sanfey, 2009).

explicitly based on dual-process ideas. Hence, the model is neither a multiple selves one nor an “as if” construction. Rather, the objective of the exercise is to show that formal modeling of dual-process theories (which is essentially absent in psychology) can result in interesting insights while remaining grounded on empirical evidence on human decision processes. The hope is that the model will provide a methodological bridge between economic modeling and actual (and mainly empirical) psychological research.

The basic ingredients of the model are as follows. The general framework starts by distinguishing two broad types of processes, controlled/deliberative and automatic/impulsive. In contrast to multiple-selves models, we propose to explicitly incorporate a differential modeling for both kinds of processes. Controlled processes will be associated to myopically rational agents who optimize a given utility function under correctly updated beliefs. Even more, we will consider a single, fixed utility function for all controlled processes, i.e. controlled processes are always mutually consistent. As a model for automatic processes, we consider an explicit stimulus-response approach. More specifically, automatic processes correspond to the behavior of boundedly rational behavioral rules after an implicit dynamic process has “settled”. That is, an automatic process acts as if in the steady state of an appropriately defined dynamical process which itself models the evolution of behavioral rules of thumb. In this paper, we will focus on *reinforcement* (or *habit formation*) behavior, one of the basic motors of human learning (Bush and Mosteller, 1951, 1955), which was first incorporated into economic modeling by Cross (1973, 1983) and has been often analyzed in the context of economic decision making (see e.g. Charness and Levin, 2005; Achtziger and Alós-Ferrer, 2014) and incorporated in models of learning in games (Börgers and Sarin, 1997; Erev and Roth, 1998; Camerer and Ho, 1999).⁴

We will build the analysis in two steps. In the first step, we consider an isolated instance of a self-control problem, simplified to the essence: an opportunity to start a desirable behavior and a later situation where temptation appears. At each possible decision point, as postulated by (a simplified version of) dual-process theories, the actual decision of a decision maker might be made by either a controlled or an automatic process, according to probabilities exogenous to the decision maker. The model will be analyzed using the concept of *Dual-Process Equilibrium*, which takes into account the utility-maximizing character of controlled processes and the steady-state nature of automatic ones.

The exercise of willpower can be identified with the inhibition of automatic behavior (impulses) to enable controlled processes to act.⁵ Hence, the probability of inhibition

⁴There is abundant psychophysiological evidence showing that reinforcement processes have an automatic nature, being associated with extremely fast and unconscious brain responses. See e.g. Holroyd and Coles (2002).

⁵The detection of a conflict between competing processes is one of the functions of the anterior cingulate

is the essence of self-control. Many factors can influence this personal parameter, from individual characteristics to training or the nature of the problem. In the second step, we will introduce a “personal evolution” approach, internalizing this probability. To keep the analysis simple, personal evolution is identified with a local-adjustment dynamics where the probability increases or decreases following the gradient of individual payoffs. Implicitly, the decision maker learns to adjust the equilibrium of the various processes determining behavior in a way which is rationally responsive to the anticipated results. However, this leaves the door open to myopic behavior in the form of local maxima. Hence, even in the personal evolution extension of the model, there is no all-knowing, long-run planner.

We find that decision makers can become trapped in suboptimal situations with frequent self-control failures. Depending on situational and personal factors to be described below, those can be of different kinds. It is possible that a decision maker becomes trapped in an “excessive deliberation” state, where he consumes his cognitive resources deliberating on whether to get started, but frequently fails to do so, even though a lower reliance on deliberation would result in a higher rate of success.⁶ Decision makers in this situation will not face temptation often, because their main problem is a (rationalized) failure to get started. On the opposite side of the behavioral spectrum, it is possible that deliberation always leads to getting started, but still self-control problems appear. One problem is that a decision maker might end up relying excessively on automatic processes, even though more frequent deliberation would be optimal. A second and more surprising problem is a mirror-image of the first. In this case, decision makers might be caught in an “excessive self-control” trap, where they frequently exercise effortful self-control, and then end up giving in to temptation too often after having started a goal-directed plan of action. In both cases, decision makers will identify self-control failures with an inability to resist temptation or shield already-initiated goal-striving.

As commented above, the model at hand draws from the dual-process literature from psychology and is only conceptually related to multiple selves models. The most closely related papers in the economics literature are Bernheim and Rangel (2004) and Benhabib and Bisin (2005). Bernheim and Rangel (2004) propose a model for the study of addiction and postulate a “hot” and a “cold” mode of decision. The hot mode corresponds to an

cortex (ACC), a part of the brain located in the medial part of the frontal lobes (Botvinick, Cohen and Carter, 2004; Ridderinkhof et al., 2004). Activity in the ACC has been related to conflict detection and inhibition of automatic responses in a large number of tasks, including decision-making under uncertainty (De Neys, Vartanian and Goel, 2008; Achtziger et al., 2014).

⁶An additional negative effect of excessive deliberation pointed out in psychology is that it might also lead to inaction and procrastination simply because a decision is not reached in time. We abstract from this added complication here.

automatic, non-rationalizable process which always leads to drug usage, while the cold mode reflects rational, forward-looking considerations. This model, as the one proposed here, is closer to dual-process theories in that the automatic process is not endowed with a utility function but rather used to directly model a psychological impulse (drug consumption). The main difference is that Bernheim and Rangel (2004) aim to model long-run management of addictive tendencies and hence their controlled process is hyperrational in the sense of long-run intertemporal optimization given a state space reflecting levels of addiction. A second difference is that, while our controlled processes are less sophisticated than those of Bernheim and Rangel (2004), our automatic processes are more sophisticated than theirs, since they will not prescribe a unique mode of behavior, but rather reflect the probabilistic outcome of a reinforcement dynamical process. At a conceptual level, Bernheim and Rangel (2004) concentrate on a behavioral problem which dominates the lives of affected agents, and their modeling choice is appropriate to study the long-run struggle against addiction. In contrast, we concentrate on day-to-day self-control behavior in less dramatic situations. Benhabib and Bisin (2005) consider a consumption-savings model where an automatic process leads to overconsumption and needs to be inhibited. In case of inhibition, a controlled process takes over and determines consumption by optimizing the intertemporal consumption path. The automatic process can be rationalized via utility maximization under incorrect assumptions (see Benhabib and Bisin, 2005, Appendix A), although this is not an essential feature of the model. As in the case with Bernheim and Rangel (2004), the main difference with our approach is that our controlled processes consider a local, one-shot situation and do not solve intertemporal optimization problems. Again, the modeling decision of Benhabib and Bisin (2005) is appropriate in a setting with intertemporal constraints, where the object of interest is precisely the distribution of consumption over time, while we target simple self-control problems which are faced repeatedly.

The paper is structured as follows. Section 2 discusses evidence on self-control problems and presents the basic model. Section 3 presents the concept of Dual Process Equilibrium and characterizes the equilibria of the model. Section 4 tackles the personal evolution approach. Section 5 briefly discusses alternative approaches, including sophisticated long-run selves and more biological evolutionary approaches. Section 6 concludes. Proofs are relegated to the Appendix.

2 A Self-Control Model

In order to illustrate the approach, we will focus on *simple self-control problems*, of the kind encountered by many decision makers in their daily lives. Such problems are characterized by being *frequent*, *mild*, and *local*. By *frequent* we mean problems which decision makers face often and are recurrent in their personal experiences. For instance, sitting down to write one of the referee reports in the pile and then actually completing it within a short amount of time is a frequent task for many researchers. Likewise, completing the daily exercise routine or making sure to avoid both French fries *and* dessert while ordering in a restaurant are problems encountered recurrently, while making sure to make an appointment for your yearly medical check up or completing your tax return forms are not. We explicitly constrain ourselves to *mild* problems, e.g. excluding addiction or other problems which come to dominate the decision-maker's life, giving rise to conscious strategizing and the recruitment of external help. Finally, by *local* we mean problems which are typically dealt with one at a time and within a relatively short time interval, whose intertemporal nature is not salient enough to demand active long-run planning. For instance, for most people avoiding the second glass of wine at dinner or getting through the afternoon without a cigarette, an extra cup of coffee or a bag of potato chips fall into this category. Likewise, in most psychologically healthy decision makers, compulsive buying in a specific domain (say, DVDs or shoes) might be perceived as a (mild) problem, but the associated level of spending will rarely make explicit the need to optimize intertemporal consumption. In contrast, saving for your pension or paying off an overwhelming debt is by definition a non-local problem, which will make salient that decisions should not be considered in isolation as they appear.

We formulate the model in a stylized way which aims to capture the essence of simple self-control problems. A decision maker (DM) faces a task which requires perseverance to complete. Initially ($t = 1$), the DM can simply *Shirk*, not getting started at all, resulting in a neutral outcome, also called *Shirk* for simplicity. If the DM actually manages to *Start* with the task, further down the road ($t = 2$) he will be faced with temptation. The DM will then have a choice between exercising self-control (*Persevere*), resulting in *Success*, or alternatively to *Give Up*, resulting in failure. A graphical representation of the decision tree is given in the left-hand part of Figure 1. The decision to Start might refer to either actually starting a task (e.g. exercising, starting the referee report or the grading of the pile of exam papers) or to the inhibition of an initial temptation (the first cigarette or cup of coffee, ordering a steak at dinner). Accordingly, the decision to Persevere might refer to either working towards task completion (keep running after the first few miles, complete the report after the first interruption, keep grading exams) or to the inhibition

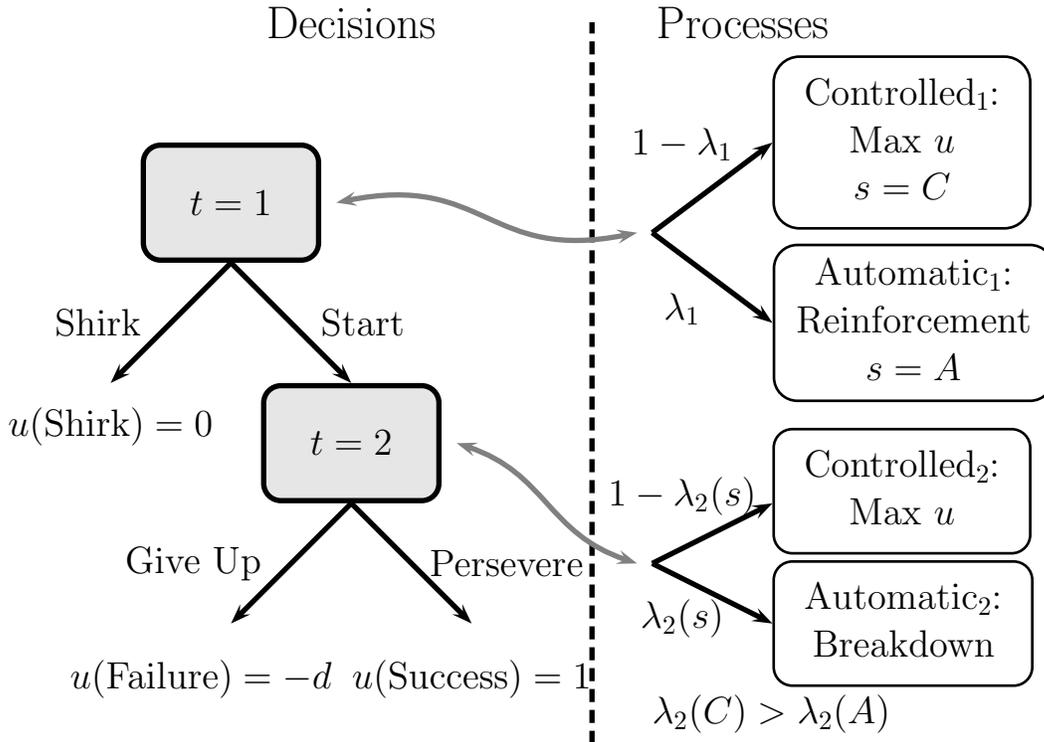


Figure 1: Schematic representation of the basic model.

of a further onset of temptation (not getting up to get a coffee or the bag of potato chips later in the afternoon, not ordering desert when the waiter asks).

This setting is similar to the basic problem described in e.g. the multiple-selves model of Bénabou and Tirole (2004), who use it to study internal commitments based on self-reputation over the own, uncertain willpower. Here, we postulate a model explicitly based on dual-process ideas instead. To avoid confusion, we will distinguish between *self-control* and *willpower*. We will use the word *self-control* in a general sense referring to all self-regulation strategies that a decision-maker might use, including e.g. adopting self-commitment mechanisms, psychological interventions as the implementation intentions mentioned above, or the avoidance of future temptations (Thaler and Shefrin, 1981). The word *willpower* will be used in a more specific sense, namely the effortful inhibition of impulses as part of a self-control act.⁷ In particular, instances of self-control without conscious deliberation (e.g. Bargh and Chartrand, 1999) correspond to effortless self-

⁷Self-control is often defined in the second, narrower sense (e.g. Baumeister, 2002). Recently, Fujita (2011) has argued in favor of a broader conceptualization, defining self-control as the process of advancing distal, abstract motives over more proximal, concrete ones when the two motives conflict. Under this view, for instance, the deliberate use of a self-commitment mechanism is a self-control strategy, whether the mechanism is external (as e.g. committing to savings plans) or internal (as in e.g. Ali, 2011).

control, i.e. self-control not requiring willpower. The existence of this possibility is the key to the success of implementation intentions.

Psychological models of self-control (Baumeister, 2002) postulate that the same mental resource is used for different tasks requiring willpower, ranging from the inhibition of impulses to persistence on complex cognitive tasks, etc. This resource is viewed as being limited; acts of willpower consume it. In other words, decision makers are on a “cognitive budget” regarding willpower. Limiting this resource by exerting effortful self-control induces a state of “ego-depletion” that impairs the performance of subsequent tasks that also require willpower. Accordingly, ego-depletion is a state in which decision makers’ capacity to exert effortful self-control is reduced (Baumeister, 2002; Baumeister, Vohs and Tice, 2007; Muraven, Tice and Baumeister, 1998).⁸ Dozens of empirical studies have found that depleted subjects have more difficulties exercising willpower than non-depleted controls, and are hence more likely to give up in tasks requiring persistence (Baumeister et al., 1998), overeat (Vohs and Heatherton, 2000), shop impulsively (Vohs and Faber, 2007), act selfishly in the Dictator Game (Achtziger, Alós-Ferrer and Wagner, 2014), and even cheat when reporting their own performance (Mead et al., 2009). Hagger et al. (2010) present an extensive review of experimental studies in this area.⁹

We integrate these insights in our model as follows. At each of the two decision nodes, two different processes compete for determining the DM’s actual decision, an automatic process and a controlled one. Which process actually makes the decision at each node is determined probabilistically in a way which incorporates previous depletion. Further, the decision arising from each process might itself be non-deterministic. A schematic representation of the postulated processes and their probabilities is given in the right-hand part of Figure 1.

Controlled processes at $t = 1$ and $t = 2$. Controlled processes in both decision nodes correspond to utility maximization given the appropriate beliefs (which in equilibrium will be further required to be correct, given the behavior prescribed by other processes). Given preferences $Success \succ Shirk \succ Failure$, we can normalize utilities to

⁸This is often called “the strength model of self-control”. This model views self-control analogously to a muscle: behavior that requires high willpower consumes strength and energy. Willpower becomes “fatigued” after being exerted in one task, leading to a reduced self-control capacity in a subsequent task. Some studies argued that this resource could have a physical nature and be related to brain glucose (Gailliot et al., 2007; Masicampo and Baumeister, 2008). However, more recent studies present evidence against this hypothesis (Molden et al., 2012) and argue in favor of a purely motivational view (Inzlicht and Schmeichel, 2012).

⁹Ozdenoren, Salant and Silverman (2012) consider a formal model where an agent with limited willpower has to allocate this resource optimally in order to regulate consumption over time. Fudenberg and Levine (2012) modify the dual-self model of Fudenberg and Levine (2006) allowing control costs to depend on the amount of willpower used in the recent past.

$u(\text{Success}) = 1$, $u(\text{Shirk}) = 0$, and $u(\text{Failure}) = -d$ with $d > 0$.

Note that controlled processes maximize the same, shared utility function and hence are consistent with each other. This is in stark contrast to models of multiple-selves or time-inconsistent preferences, where decisions taken at different decision nodes by otherwise rational selves follow different utility functions.

Automatic process at $t = 1$. The automatic processes which play a role at each decision node capture impulsive behavioral factors. At the first decision node ($t = 1$), the automatic process reflects habit formation, determining behavior on the basis of associations created in previous instances of the problem. That is, this process is not deterministic, but rather prescribes to start or shirk with a certain probability arising from an effortless *reinforcement* process.¹⁰ Hence, while we can adopt a static approach for the controlled processes, modeling this automatic process requires an implicitly dynamic approach. Formally, this process specifies a probability to start, z_n , which is increased each time that the two-stage decision situation is faced and a decision to start (independently of which process determined it) results in success. Analogously, z_n is decreased when such a decision results in a failure. The subindex n refers to occurrences of the two-stage problem, i.e. to an orthogonal temporal dimension. In equilibrium, and following a reduced-form dynamic formulation, the process will simply be required to have arrived at a (stable) steady state z^* , i.e. $E(z_{n+1}|z_n = z^*) = z^*$, given the behavior prescribed by other processes.

Remark 1. One should be careful with the interpretation of an “automatic start”. Of course, nobody spontaneously finds himself out in the street jogging right after waking up (which is quite fortunate in the case of cold-climate countries). However, action psychology indicates that, even if the implementation is far from instantaneous, the decision itself (even the decision to go jogging in the morning) can be automatized (e.g. Achtziger and Gollwitzer, 2010; Bayer et al., 2009). Indeed, this possibility is the key for the success of psychological strategies as implementation intentions.

Automatic process at $t = 2$. At the second decision node ($t = 2$), the automatic process is meant to capture a completely different phenomenon, namely “breakdown of will” due to instantaneous temptation, or, in other words, the willpower failure which leads to not completing the task after having initiated it. The simplest way to model this phenomenon is to assume that this automatic process simply leads to giving up with certainty. Of course, this could be rationalized by an alternative utility function or

¹⁰Contrary to intuition, automatic processes are not considered deterministic in psychology. Reinforcement, which is typically accepted as one of the basic building blocks of automatic behavior, is inherently a random process which allows to tailor behavior to experience.

generalized to a probability distribution. This simple formulation is meant to capture the possible failure of self-control at this point without introducing unnecessary complications.

We want to emphasize that the phenomenon captured by the automatic process at $t = 2$ is different from automatic behavior at $t = 1$. From the point of view of rational or rationalizable behavior, there can be no possible reason to give up at $t = 2$, and yet this is a common occurrence in actual human experience. A possible way to interpret breakdown of will from a rationalistic point of view is the possibility that short-term utility (temptation) temporarily overcomes any other consideration.¹¹

The last element of the model specifies *which* process actually determines the decision at each decision node. We adopt a probabilistic formulation. Let λ_t denote the probability that the decision at node $t = 1, 2$ is determined by the automatic process. That is, $1 - \lambda_t$ is the probability of exercising willpower. Hence, a large λ_1 indicates low self-control at $t = 1$ (see again the right-hand part of Figure 1). In this framework, willpower is the act of inhibiting an automatic decision and adopting a controlled one instead. As discussed above, exercising willpower is costly and reduces the self-control budget available for subsequent decisions. This is modeled as follows. While λ_1 is fixed at $t = 1$, there are two possible values of λ_2 , which depend on whether willpower has been already exercised or not. More specifically, if the automatic process took the decision at $t = 1$, then $\lambda_2 = \lambda_2(A)$; if it was the controlled process which took the decision at $t = 1$, then $\lambda_2 = \lambda_2(C) > \lambda_2(A)$. That is, the exercise of willpower at $t = 1$, resulting in a controlled decision, reduces the probability that the DM will be able to exercise self-control again at $t = 2$. Note that an automatic decision might result in an “automatic start” with probability z^* (and hence be viewed as effortless self-control) and still consume no self-control resources. Symmetrically, a controlled decision might lead to a rational failure to start and still consume those resources. What matters is the adoption of a controlled decision, and not the particular content of that decision; to put it simply, effort is spent whenever willpower is exercised, independently of success.

¹¹A natural generalization would be to allow the automatic process at $t = 2$ to randomize between Give Up and Persevere, e.g. as the result of a random temptation utility realized according to a fixed distribution. Since the controlled process at $t = 2$ always perseveres, it is easy to see this that the model would be unchanged with a reinterpretation of the parameters, the probability of automatic perseverance being now subsumed in the probability of a decision by the controlled process. Further, suppose the decision of the automatic process at $t = 2$ would be probabilistic with a probability adjusted according to a reinforcement process as in $t = 1$. This would not capture the phenomenon we are interested in. The reason is that, if reinforcement would be based on the actual utility of persevering or giving up, as opposed to temptation, and since the result of persevering at $t = 2$ is always a success, the process would converge to the automatic process always persevering. This is, of course, a rather unrealistic result.

3 Dual-Process Equilibria

The key concept for the analysis of dual-process models as sketched here is that of a *Dual-Process Equilibrium*, which is defined to be a profile formed by optimal strategies for the controlled processes and steady states for the automatic processes. In general, optimality of the strategies of the former might depend on the actions of the latter, and the steady-state computation for the latter might depend on the strategies of the former. The key difference between this equilibrium concept and, say, a Perfect Bayesian Equilibrium, is that automatic processes are not rationalized as optimizing agents, but rather they are required to act as in a stable state of an implicit dynamical process, taking the behavior of controlled processes as given. In contrast, in a Dual-Process Equilibrium controlled processes behave in a more “classical” way, essentially corresponding to standard, optimizing players who take the behavior of other processes (controlled or automatic) fully into account.

Let us turn to the self-control problem described above. Behavior at $t = 2$ can be characterized in a straightforward way. Since $u(\text{success}) > u(\text{failure})$, the controlled process always chooses to Persevere, while the automatic process, by construction, always chooses to Give Up. The equilibrium behavior of the controlled process at $t = 1$ corresponds to *sequential rationality*, i.e. utility maximization given correct beliefs.¹² The computation is as follows.

$$(1) \quad \text{Start} \succeq \text{Shirk} \iff 1 - \lambda_2(C) - d\lambda_2(C) \geq 0 \iff \lambda_2(C) \leq \frac{1}{1+d}$$

Last, we discuss the behavior prescribed by the reinforcement process at $t = 1$. Suppose that the two-stage problem is faced repeatedly, instances being given by $n = 1, 2, \dots$. The reinforcement process determines whether to start or not following a simple probability distribution $(z_n, 1 - z_n)$. The probability of starting is updated every time this action is adopted: it is increased if it previously led to a success and decreased if it previously led to a failure. Following a classic formulation (Cross, 1973),

$$z_{n+1} = \begin{cases} \theta + (1 - \theta)z_n & \text{if success} \\ (1 - \theta)z_n & \text{if failure} \end{cases}$$

where $0 < \theta < 1$ is an adjustment factor (which will be irrelevant for steady-state compu-

¹²An extensive form capturing the model would include a move by nature after the controlled process decides to start, with fixed probabilities $\lambda_2(C)$, $1 - \lambda_2(C)$ leading to (decisions themselves leading to) failure and success, respectively. Hence the implicit condition on beliefs is that the probabilities of nature’s decisions are “known” by the rational processes.

tations). If the DM shirks, the payoffs of *start* are not experienced and z_n is not updated. If the DM starts, whether the process which led to starting the task was an automatic or a controlled one is irrelevant for the purposes of updating z_n : the probability is based on actual experiences and not on internal processes.

Given the behavior prescribed by the other three processes, it is possible to compute the expected value of z_{n+1} after the n -th instance of the problem. In equilibrium we will require the reinforcement process to have reached a probabilistic steady state, i.e. $z = z^*$ with

$$E(z_{n+1}|z_n = z^*) = z^*.$$

Further, we will also require that the one-dimensional dynamics points towards z^* in a probabilistic sense, i.e. $E(z_{n+1}|z_n < z^*) > z_n$ and $E(z_{n+1}|z_n > z^*) < z_n$ in an interval around z^* . If these conditions are met, we say that z^* is a *stable steady state* for the reinforcement process.

In view of these considerations, and in particular since behavior at $t = 2$ is already determined, we can provide a simpler definition for the model at hand.

Definition 1. A **Dual-Process Equilibrium** (DPE) of the self-control model is a pair (x, z) such that

- (i) the strategy prescribing to start with probability x is optimal for the controlled process at $t = 1$ given $\lambda_2(C)$, and
- (ii) z is a stable steady state for the reinforcement process at $t = 1$, given the behavior of the controlled process at $t = 1$.

In general, one might want to allow the increase in λ_2 due to an initial exercise of willpower to depend on λ_1 , the probability of impulsive decision-making at the beginning. For instance, it would be sensible to assume that increased self-control possibilities at $t = 1$ (reduced λ_1) result in a larger loss in the self-control budget at $t = 2$ (increased $\lambda_2(C) - \lambda_2(A)$). For this reason, and also in order to minimize the use of subindices in the subsequent analysis, we now adapt our notation as follows.

$$\lambda_1 = \Delta, \quad \lambda_2(A) = \lambda, \quad \lambda_2(C) = \lambda + g(\Delta)$$

with $g(\cdot)$ a nondecreasing function of Δ such that $0 < g(\Delta) < 1 - \lambda$. Note that we do not exclude the possibility of a constant $g(\Delta)$.

With this reformulation, the parameter λ can be interpreted as the intensity of the temptation at $t = 2$, while Δ (or, rather, $1 - \Delta$) can be interpreted as the intensity with

which willpower is exercised. From now on, we assume that $1 - \lambda(1 + d) > 0$, or

$$(2) \quad \lambda < \frac{1}{1 + d}.$$

This condition serves to exclude trivial cases. If it were violated, then the controlled process would never consider starting, even in the absence of depletion effects.

Remark 2. The optimal decision of the controlled process in a DPE, as captured by (1), does not depend on the steady-state probability of an automatic start z . This analytically convenient fact is natural in this setting. The controlled process only considers a decision conditional on having been selected to actually make that decision (i.e. given that willpower has been exercised), while the probability of an automatic start is only relevant in the counterfactual situation where the automatic process has been selected.

3.1 The Reinforcement Process

Consider a Dual-Process Equilibrium where the controlled process at $t = 1$ starts with probability $x \in [0, 1]$. This probability will affect the probability of an automatic start, z , since a controlled start leads to a decision at $t = 2$ with reduced probability of perseverance.

Consider the n -th instance of the two-stage problem. If the automatic process takes the decision at $t = 1$, then the expected value of z_{n+1} is given by

$$(3) \quad z_n [(1 - \lambda)(\theta + (1 - \theta)z_n) + \lambda(1 - \theta)z_n] + (1 - z_n)z_n = z_n [1 + \theta(1 - \lambda - z_n)]$$

If the controlled process decides instead, the expected value of z_{n+1} becomes

$$(4) \quad x [(1 - \lambda - g(\Delta))(\theta + (1 - \theta)z_n) + (\lambda + g(\Delta))(1 - \theta)z_n] + (1 - x)z_n = \\ = z_n + x\theta [1 - \lambda - g(\Delta) - z_n]$$

Recall that if either the automatic or the controlled process determine the decision but lead to shirking, z_{n+1} is not updated. We obtain

$$(5) \quad E(z_{n+1}) = \Delta z_n [1 + \theta(1 - \lambda - z_n)] + (1 - \Delta) [z_n + x\theta(1 - \lambda - g(\Delta) - z_n)] = \\ = z_n + \theta [(1 - \Delta)x(1 - \lambda - g(\Delta)) - z_n((1 - \Delta)x - \Delta(1 - \lambda - z_n))]$$

The condition for a (stable) steady state of the reinforcement process becomes

$$E(z_{n+1}) \begin{matrix} \geq \\ \leq \end{matrix} z_n \iff \Delta z_n^2 + z_n ((1 - \Delta)x - \Delta(1 - \lambda)) - (1 - \Delta)x(1 - \lambda - g(\Delta)) \begin{matrix} \leq \\ \geq \end{matrix} 0$$

The left-hand-side of the inequality is a polynomial of second degree with positive discriminant, hence its value is positive below the first root and above the second one, and negative between the roots. Thus z_n increases in expected value between the roots and decreases otherwise. However, the value of the polynomial at $z_n = 0$ is negative and can be zero only if $x = 0$. Hence the first root is either strictly negative and can be ignored (if $x > 0$) or it yields a steady state $z = 0$ which is then unstable (if $x = 0$). The value of the polynomial at $z = 1$ is strictly positive for any value of x . Hence, it follows that there exists a unique steady state, which is (globally) stable and is (after some simplifications) given by

$$(6) \quad z^* = \frac{1}{2\Delta} \left[-((1 - \Delta)x - \Delta(1 - \lambda)) + \sqrt{((1 - \Delta)x + \Delta(1 - \lambda))^2 - 4x\Delta(1 - \Delta)g(\Delta)} \right]$$

3.2 Equilibria

We start by analyzing pure strategy equilibria¹³ where the controlled process at $t = 1$ always shirks, i.e. $x = 0$. By condition (1), this occurs if and only if

$$\lambda + g(\Delta) \geq \frac{1}{1 + d}.$$

If the controlled process always shirks, equation (5) reduces to

$$(7) \quad E(z_{n+1}) = z_n + \Delta\theta z_n (1 - \lambda - z_n)$$

and there are two steady states, $z = 0$ and $z = 1 - \lambda$. The first one is unstable and the second one is stable. Hence, we obtain a Shirking DPE where the controlled process at $t = 1$ always shirks and the reinforcement process starts automatically with probability $z = 1 - \lambda$.

The interpretation of this equilibrium is straightforward. The controlled process is discouraged by a too large probability of breakdown of will at $t = 2$, i.e. whenever the DM “thinks about it”, he gives up since he does not trust himself to continue. Given this, the action to start occurs only automatically (“just do it”), and the model becomes a classic reinforcement example. Hence, in equilibrium the probability to start is numerically

¹³Whether a Dual-Process Equilibrium is in mixed or pure strategies refers to the strategies of the controlled processes only.

equal to the probability that the task is successfully completed, corresponding to a simple *probability matching* result as known from the reinforcement learning literature (Bush and Mosteller, 1951; Siegel and Goldstein, 1959).¹⁴

We now turn to pure strategy equilibria where the controlled process at $t = 1$ always chooses to start if willpower is exercised, i.e. $x = 1$. By condition (1), this occurs if and only if

$$\lambda + g(\Delta) \leq \frac{1}{1+d}.$$

If the controlled process always chooses to start, equation (5) reduces to

$$(8) \quad E(z_{n+1}) = z_n + \theta [(1 - \Delta)(1 - \lambda - g(\Delta)) - z_n(1 - 2\Delta + \Delta\lambda + \Delta z_n)]$$

There is a unique steady state, which is globally stable and given by (taking $x = 1$ in (6))

$$(9) \quad z^* = 1 + \frac{1}{2\Delta} \left[-(1 + \Delta\lambda) + \sqrt{(1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g(\Delta)} \right]$$

Note that, even though in this kind of equilibria the controlled process always chooses to start, this does not mean that the automatic process does not influence behavior, since whether one or the other process actually takes the decision depends on whether willpower is exercised.

Finally, consider mixed-strategy equilibria where the controlled process at $t = 1$ randomizes, starting with probability $0 < x < 1$. By condition (1), this can occur if and only if

$$\lambda + g(\Delta) = \frac{1}{1+d}.$$

Since this equation involves only exogenous parameters at this point, we conclude that mixed-strategy equilibria are possible but correspond to knife-edge situations. If the equality holds, we obtain a continuum of mixed-strategy equilibria with $x \in [0, 1]$ and $z = z(x)$ given by (6).

In summary,

Theorem 1. *A Dual-Process Equilibrium of the self-control model always exists.*

- (a) *If $\lambda + g(\Delta) > \frac{1}{1+d}$, there exists a unique DPE, where the controlled process at $t = 1$ always shirks ($x = 0$) and the reinforcement process starts automatically with probability $z = 1 - \lambda$ (Shirking DPE).*

¹⁴Suppose a decision maker repeatedly chooses among two options, A and B, only one of them leading to a reward. Suppose that the probability that A is the winning option is constant among periods and equal to $p > 1/2$. In the long run, reinforcement learning will lead to the decision maker choosing A a percentage p of the time, while a rational decision maker should simply choose A all the time.

- (b) If $\lambda + g(\Delta) < \frac{1}{1+d}$, there exists a unique DPE, where the controlled process at $t = 1$ always starts ($x = 1$) and the reinforcement process starts automatically with probability z^* given by (9) (Willpower DPE).
- (c) If $\lambda + g(\Delta) = \frac{1}{1+d}$, there exists a continuum of DPE, indexed by $x \in [0, 1]$. At each one of these DPE, the controlled process at $t = 1$ starts with probability x and the reinforcement process starts automatically with probability z^* given by (6).

3.3 Rational Failures to Get Started

The existence of Shirking equilibria in Theorem 1 captures the phenomenon of “deliberative failures to get started”. Case (a) indicates that this phenomenon might appear when depletion effects are strong or the disutility of failure is large. For instance, imagine a student who suddenly realizes that there are only a few weeks left until the exam for a difficult, boring subject (and let us abstract from the reasons for the student’s procrastination). In order to pass the exam, the student needs to tackle an intensive study plan with several multi-hour study sessions per day. Every given day, it is likely that the first session will cause a strong depletion effect leading to a failure of perseverance as the time slot for the second session approaches. The student might well anticipate this and never get started at all, giving up on the subject. This is related to the *sophistication effect* described by O’Donoghue and Rabin (1999) in the context of a quasi-hyperbolic discounting model. If an agent is fully aware of future self-control problems, this awareness influences present behavior.¹⁵ A different example are tasks where there is a high social or subjective cost of failure, e.g. if giving up carries a (subjective or objective) stigma. For instance, starting with but failing to complete an ambitious task might be damaging for the DM’s self-image, and hence the DM might rationally decide not to risk failure and never attempt the task.

A shirking equilibrium, however, is an equilibrium, and under the condition given in Theorem 1(a), it is the only equilibrium. It would be futile to attempt to bring the DM to exercise willpower and deliberately start the self-control task. The effects of depletion, which the DM anticipates, will offset the benefits of deliberately starting more often. Depletion will lead to increased failure (and hence increase disutility). Increased failure will feed back to the automatic process, which will correspondingly reduce the share of automatic starts.

¹⁵O’Donoghue and Rabin (1999, 2001) capture this kind of observations with interpersonal comparisons between *sophisticate* agents and *naïve* ones. In contrast, the model at hand relies on intrapersonal heterogeneity. The condition in Theorem 1(a) specifies when the sophistication effect will take over for a particular decision maker.

This outcome is, however, suboptimal. In a shirking equilibrium, the automatic process is caught in an “irrational” probability-matching outcome. If depletion effects were absent, under condition (2) it would be optimal to start at $t = 1$. Hence, it is clear that the payoffs of a shirking equilibrium could be increased if the automatic process could be brought to start more often, without incurring in depletion effects. More formally, the expected utility in shirking equilibria (where $x = 0$ and $z = 1 - \lambda$) is

$$(10) \quad \Pi_0 = \Delta(1 - \lambda)(1 - \lambda(1 + d)) > 0$$

and, if the automatic process would operate with a start probability of $z > 1 - \lambda$, the expected payoff would be $\Delta z(1 - \lambda(1 + d)) > \Pi_0$. Moreover, increasing Δ (thinking less) would also result in an increased utility, even if z remained unchanged.

Increasing Δz without engaging active deliberation is precisely the objective of psychological manipulations to counteract the failure to get started, as e.g. implementation intentions. By setting a just-do-it automatized action trigger, an implementation intention simultaneously increases the likelihood of automatic behavior (reducing willpower exertion) and the probability that the automatic process starts with the task. This explains why, in situations where we can rationally talk ourselves out of an action plan, a seemingly irrational commitment to give up willpower acts might be welfare-improving.

4 Personal Evolution and Optimal Willpower

Many decision makers struggle for self-control in certain aspects of their lives. Here we propose a “personal evolution” approach as a shortcut for *intrapersonal* parameter adjustment, as (over the course of one’s life) the reliance on certain processes becomes established.

More specifically, suppose the decision maker faces simple self-control problems as described above along his or her personal life. That is, the decision maker is frequently and recurrently confronted with mild self-control problems which he deals with on a local, one-at-a-time basis, from whether or not to finish the bag of potato chips or drink the second glass of wine to whether or not to go out jogging or finish the grading of the pile of exam papers. In spite of the frequency of such problems, none of them are life-threatening or, taken individually, even deserving of too much attention. The quest for optimal self-control reduces to adjustments in the parameter Δ , the probability of impulsive behavior in the presence of full self-control resources. Imagine that the decision maker (consciously or not) attempts to slowly build self-control (reducing Δ) but repeatedly experiences

failure, in the form of giving in to temptation, after exercising willpower. Given the disutility of failure, this might result in an overall reduction in utility, discouraging the decision maker from further reducing Δ . Reciprocally, if the reduction in Δ leads to an overall increase in utility due to an increase in the success rate, the decision maker will further reduce Δ .

In this section, we internalize the parameter Δ following a further implicitly dynamic approach. Our interpretation is that, for classes of decision problems which are encountered relatively often, this parameter is “adjusted” over one’s life in a local way, i.e. it will increase or decrease slightly whenever doing so results in a utility gain. In other words, we adopt a gradient dynamics leading to *local* optimality as a criteria for personal optimality. In the framework of the self-control model analyzed in the previous sections, we then ask the question of whether relying on automatic processes beyond a given, unavoidable background level (presumably founded on biological evolution, childhood experiences, etc) can become a local optimum for a decision maker.

Consider the following framework. Let $g(\Delta) = g$, i.e. there is a fixed, physiological depletion effect. The utilities of shirking, success, and failure are fixed. The parameter λ is also fixed and, as commented above, taken to represent a measure of the “temptation” associated to a given self-control problem. However, the parameter Δ might vary as follows:

$$\Delta = \Delta_0 + \Delta_p$$

where Δ_0 is a fixed “background automaticity rate” (you cannot avoid a certain probability of reacting automatically), but Δ_p is a “personal” parameter which becomes trained over time. An appropriate criterion to evaluate “personal optimality” is simply expected utility.

$$(11) \quad \begin{aligned} \Pi(x, z) = \Delta z [-\lambda d + (1 - \lambda)] + (1 - \Delta)x [-(\lambda + g)d + (1 - \lambda - g)] = \\ \Delta z [1 - \lambda(1 + d)] + (1 - \Delta)x [1 - (\lambda + g)(1 + d)] \end{aligned}$$

Using expected utility as welfare criterion on the (x, z) space is equivalent to taking the point of view of a benevolent planner who examines the short-term utility obtained by a decision maker whose personal parameters are fixed. That is, the hypothetical planner computes the utility of a decision maker in full knowledge that, as long as x and z have interior values, the decision maker will sometimes shirk, and sometimes start and then give up. Since in equilibrium x and z are determined by the decision maker’s personal parameters, we can analyze questions of optimality and personal evolution based on short-term payoffs on this space. We will briefly comment on some alternative approaches in Section 5.

We focus on the case where Willpower equilibria exist, since else there is little to study. The formalism is particularly convenient here, since for this particular problem the parameter Δ does not affect whether the DPE is of one or the other type (this needs not be true for more general models).

The utility in Willpower equilibria, where $x = 1$ and $z = z^*$, is given by

$$\Pi_1 = \Pi(1, z^*) = \Delta z^* [1 - \lambda(1 + d)] + (1 - \Delta) [1 - (\lambda + g)(1 + d)]$$

which, substituting z^* as given by (9)) and simplifying, yields

$$(12) \quad \Pi_1 = (1 - \lambda(1 + d)) \left[1 + \frac{1}{2} \left[-(1 + \Delta\lambda) + \sqrt{(1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g} \right] - (1 - \Delta) \left(\frac{g(1 + d)}{(1 - \lambda(1 + d))} \right) \right]$$

Say that Δ is a *personal optimum* if it is a local maximum of Π_1 under the constraint that $\Delta \geq \Delta_0$. Say that it is *efficient* if it is a global optimum of Π_1 under the same constraint. The following lemmata (whose proofs are relegated to the appendix) explore the properties of Π_1 .

Lemma 1. Π_1 is a strictly convex function of Δ .

Lemma 2. There exist \underline{g}, \bar{g} with $0 < \underline{g} < \bar{g} < 1$ such that

- (a) if $g \geq \bar{g}$, Π_1 is strictly increasing in Δ ;
- (b) if $g \leq \underline{g}$, Π_1 is strictly decreasing in Δ ; and
- (c) if $\underline{g} < g < \bar{g}$, Π_1 has a unique interior minimum $\Delta(g)$.

The convex shape of Π_1 has an intuitive interpretation. Consider case (c) in the last lemma. The utility function Π_1 has two local maxima, with minimal and maximal willpower, respectively. The reason for this U-shaped utility is the tradeoff between the direct and indirect effects of willpower exercise. Increasing willpower (reducing Δ_p) leads to a larger share of deliberate starts. However, each deliberate start causes depletion and reduces the chances of later perseverance. Further, failed perseverance “teaches” the automatic process (through reinforcement) to reduce the probability of an automatic start.

Suppose that the DM is exercising very low willpower (Δ close to 1). Then, most of the successful runs (starting, then persevering) come from automatic starts, which occur with a given probability $z > 0$. If the percentage of the decisions taken by the controlled

process is slightly increased, the associated successes will deliver positive utility, but the increased failures due to the effects of depletion will feed back in an overall reduced share of automatic starts. Since most of the successful runs actually come from automatic starts, this will have a large negative effect on overall utility.

Suppose, on the contrary, that the share of automatic decisions is small (Δ close to 0). Then, the positive effect of increasing the number of deliberative starts will easily offset the negative effect of reduced automatic starts due to failed attempts. However, if depletion effects are too strong, as in case (a), the direct disutility of failure is too large and this offsetting effect will never materialize. On the other hand, if the effects of depletion are small enough, as in case (b), increasing the share of deliberate starts will always pay off.

Since there are at most two personal optima, efficiency becomes a matter of pairwise comparison. We conclude that, depending on personal and situational factors (the strength of the temptation, the personal disutility of failure, the background automaticity, and the depletion associated to the task), either maximum exercise of willpower or full automatization of behavior will be efficient. However, the DM might well be trapped in the inefficient personal optimum. The following theorem spells out the circumstances under which each personal optimum is efficient.

Theorem 2. *Let $\lambda + g < \frac{1}{1+d}$. If the probability of automatic decisions at $t = 1$, Δ , is constrained by $\Delta \geq \Delta_0 > 0$ then there exist $\underline{g}, \bar{g}, g^*$ with $0 < \underline{g} < g^* < \bar{g} < 1$, and continuous functions $\Delta(g)$, $\Delta'(g)$ with $\Delta'(\underline{g}) = \Delta(\underline{g}) = 1$, $0 \leq \Delta'(g) < \Delta(g)$ on $g \in]\underline{g}, \bar{g}[$, $\Delta'(g^*) = 0$, and $\Delta(\bar{g}) = 0$, such that exactly one of the following cases occur:*

- (I) **Efficient Self-Control.** *Maximum Self-Control ($\Delta = \Delta_0$) is the only personal optimum and also efficient if and only if $g \leq \underline{g}$ (Low Depletion).*
- (II) **Efficient Automaticity.** *Full automaticity ($\Delta = 1$) is the only personal optimum and also efficient if and only if one of the following possibilities occur:*
 - (IIa) High Depletion: $g \geq \bar{g}$;
 - (IIb) High Background Automaticity: $\Delta_0 \geq \Delta(g)$ (and $\underline{g} < g < \bar{g}$).
- (III) **Self-Control Failure.** *Full automaticity is a personal inefficient optimum (with maximum control being the efficient optimum) if and only if $\underline{g} < g < g^*$ (moderately low depletion) and $\Delta_0 < \Delta'(g)$ (low background automaticity).*
- (IV) **Excessive Self-Control.** *Maximum Self-Control is a personal inefficient optimum (with full automaticity being the efficient optimum) if and only if $\underline{g} < g < \bar{g}$ (moder-*

ately low depletion) and $\Delta'(g) < \Delta_0 < \Delta(g)$ (moderate background automaticity).
 In case $g \geq g^*$, then $\Delta'(g) = 0$ and the condition $\Delta'(g) < \Delta_0$ is void.

(V) In the knife-edge case with $\Delta_0 = \Delta'(g)$ and $\underline{g} < g < g^*$, both maximum self-control and full automaticity are personal efficient optima.

The different cases in the last theorem are qualitatively illustrated in Figure 2. We can summarize the classification in words as follows. First, a task causing next to no depletion (Case I) does not pose a problem (since you know you can trust yourself in the future) and will efficiently be undertaken under maximum control. This kind of tasks will typically not be perceived as self-control problems at all.

Second, if a task causes a high depletion (but not so high that the equilibrium becomes of the shirking type), then it is optimal to automatize the associated decisions as much as possible. Note that this is different from the reasoning behind a Shirking DPE. In the situation described in case IIa, the controlled process at $t = 1$ will still find it optimal to start, *given that this process is required to make the decisions*. However, the DM (viewed as a planner) would have preferred that this process does not make the decision at all. The reason is that, for the controlled process, the costs of depletion are sunk. *Ex ante*, however, it would have been preferable to avoid them.

For intermediate depletion values, the outcome depends on the background automaticity rate. If that rate is high, then in practice a full-control personal optimum becomes unattainable (independently of whether it *would* have been efficient or not), and hence again decision makers efficiently rely on full automaticity (case IIb). For instance, if you want to get rid of an habitual behavior, but that behavior happens to be aligned with your inborn (or childhood-learned) reflexes, you will find that consciously fighting it is next to impossible.

Trouble arises systematically for intermediate depletion values and moderate or low background automaticity rates, for in this case there are always two personal optima but only one of them is efficient (except in knife-edge cases). If background automaticity is relatively low (case III), we reach a situation with which many decision makers frequently identify. In this case, although exercising maximum self-control would be efficient, there exist personal optima with full automaticity. In this case, the decision-maker is caught in a trap. Although it might be clear that exercising full self-control would be beneficial, the road to heaven is paved with very slippery stones, for, since the inefficient personal optimum is a local maximum, any small increase in self-control will lead to a *decrease* in utility. Hence, slowly increasing your exercised self-control will not lead to immediate rewards; quite the opposite.

This situation, with moderate depletion and low background automaticity, corresponds to the quintessential self-control problem where people fail to exercise willpower when they should. The insight of the model presented here is that decision-makers might still be at a personal optimum, where small increases in the exercise of willpower are counterproductive. Hence, we are led to the conclusion that the most promising course of action is a radical shift, which typically will fail in the absence of external help, e.g. in the form of a commitment mechanism or psychological interventions as those described in the introduction.

The second problematic situation is an “excessive self-control trap” (case IV). If the background automaticity is relatively high but not extreme, full automaticity is efficient, but decision makers may be caught in personal optima with maximum exercise of willpower. Ultimately, this is inefficient because those attempts remain relatively unsuccessful compared with what could be achieved under full automaticity. Decision makers exercise too much willpower, constantly trying to complete a task or bring a certain habitual behavior in line, more or less successfully. Slightly lowering their efforts is counterproductive. The irony is that spontaneous, effortless, unplanned behavior would actually result in an improvement in this case. This phenomenon might be more relevant than it appears at first glance, since it stands to reason that it will typically not be fully recognized. After all, a decision maker who has fallen in this trap seems to be doing his or her best. For instance, decision makers in this case might be perceived as “willing hard-workers” who constantly overcommit and frequently start new projects, but fail to complete most of them in time. It is also conceivable that people with chronic (but non-pathological) weight problems, or people exhibiting mild forms of compulsive behavior, might fall in this case. In both examples, the behavior in question is likely to be relatively automatized (a moderately high, probably learned background automaticity rate) and the depletion associated to every act of willpower is moderate.¹⁶

¹⁶Although addiction problems are most certainly not merely willpower problems (Bechara, Noel and Crone, 2006), this might also be the case some non-compulsive smokers fall in, hence yielding an (admittedly speculative) explanation why *unplanned* smoke-quitting attempts are twice as likely to be successful as planned ones (Ferguson et al., 2009; West and Sohal, 2006). Indeed, it appears that each year more smokers quit smoking unassisted (“cold turkey”, i.e., simply stop overnight; Chapman and MacKenzie, 2010) than by all assisted methods combined. In terms of the model, this might correspond to smokers previously caught in the excessive self-control trap (since large numbers of smokers report wanting to give up smoking) and suddenly finding out that there is a minimum-resistance path to accomplish their objectives.

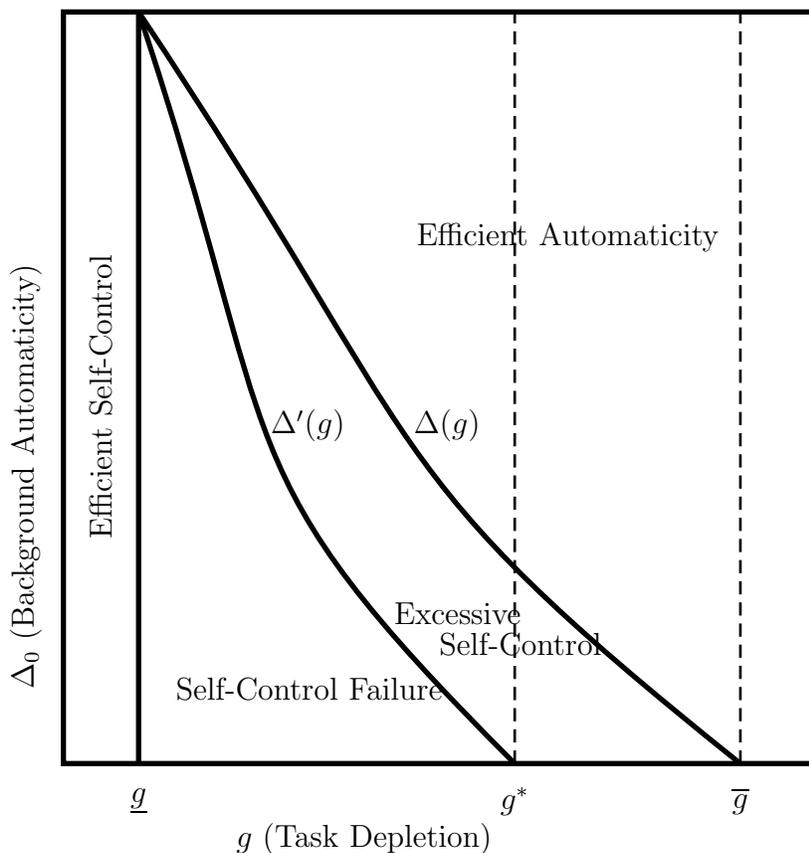


Figure 2: Illustration of the cases in Theorem 2.

5 Alternative Approaches

5.1 Sophisticated Planning and Long-Run Selves

The personal evolution approach is well suited for the kind of decisions considered here: frequent, mild, and local. For these decisions, gradual adjustment on the basis of short-term payoffs is more likely than conscious, long-run intertemporal optimization. We simply do not have the resources to maintain conscious, deliberate control over all aspects of our daily lives, and the repeated breakdowns of will most of us experience in one or other domain are a testimony to it. Self-control problems of a more severe nature, however, are more likely to take a prominent place in the decision-maker's life (e.g. intertemporal consumption and saving), and hence elicit conscious, rational planning attempts. Models along this line have been considered e.g. by Benhabib and Bisin (2005) or Fudenberg and Levine (2006) and typically include an infinitely-lived long-run planner (or long-run self) who is able to conduct intertemporal optimization but faces a sequence of short-run, impulsive selves.

The model described so far can be extended along these lines to cover problems beyond the “frequent, mild, and local” class. Essentially, the insights are similar to those outlined in the last section, and the same kind of personal traps appear. The differences concern mostly the basins of attraction of the traps. In the interest of brevity, we simply sketch the approach here.

Consider a patient long-run self facing a sequence of short-run selves. At each point in time $\tau = 0, 1, \dots$, the decision maker faces the two-step situation described in Section 2. The short-run self at time τ takes the place of the automatic process in the first subperiod. The long-run self makes all decisions for controlled processes. In addition, the long-run self has the capacity to gradually adjust the willpower parameter Δ over time, e.g. by slowly building or giving up self-control. This adjustment is costly, with the cost subtracted from instantaneous utility at t .

The first observation is that the decisions for controlled processes remain unchanged, since they are optimal and consistent with each other in the basis model, and there is no interaction with the adjustment process. The second observation is that, as Lemma 1 shows, the utility Π_1 is a strictly convex function of the willpower parameter Δ . In other words, the optimal solution is always at an extreme, $\Delta = \Delta_0$ or $\Delta = 1$.

Depending on the initial value of Δ , the long-run planner will always find it feasible to plot a path towards one of the extreme values in such a way that utilities are increasing in time. If Π_1 is decreasing at Δ , this path leads to $\Delta^* = \Delta_0$. If Π_1 is increasing at Δ , it leads to $\Delta^* = 1$. Call this the “easy path”. This path essentially mimics the predicted path under a gradient dynamics as described in the personal evolution approach. We can also consider a path where the long-run planner first heads towards the minimum, with utility decreasing along the way, and then starts climbing (with now-increasing utility) towards the other extreme. That is, this path passes through the valley of the U-shaped function Π_1 . Call it the “hard path”. As long as the agent is not too impatient and the costs of adjustment are not too large, it is easy to see that every other possibility will be dominated by one of these two paths.¹⁷

The remaining question is whether the easy or rather the hard path will be optimal. In cases (I) and (II) of Theorem 2, Π_1 has no interior minimum and hence the hard path does not exist. The decision maker will move towards the unique personal optimum, which is also efficient. In the knife-edge case (V), where both personal optima are efficient, the easy path clearly dominates the hard path and the decision maker will move towards one of the personal optima depending only on initial conditions.

¹⁷With very large impatience or very large adjustment costs, there is no scope for long-term planning and the decision maker will be stuck at his initial Δ .

In cases (III) and (IV), one of the personal optima is efficient and the other is inefficient. If the initial condition places the decision maker on the side of the efficient personal optimum, then he or she will follow the easy path there. The most interesting situation arises if the initial condition places the decision maker on the side of the personal inefficient optimum. Then, the hard path towards the efficient optimum will still be better than the easy path provided the long-run planner is patient enough, the costs of adjustment are small enough, and the size of the adjustments is not too small. The reason is that convergence towards the optimum will occur in finite time, and hence the long-run gain offsets all possible costs of adjustment and the reduced utility while “passing through the valley”, if the future is not discounted too sharply.

In summary, depending on initial conditions and parameter values as e.g. the long-run self patience, we would predict that the decision maker can end up in a personal inefficient optimum or in an efficient one. Hence both self-control traps described in the previous section are possible. The main difference concerns the basins of attraction. Under the personal evolution approach, in cases (III) and (IV) the interior minimum of Π_1 fully characterizes the boundary of the basins of attraction. Depending on which side of this boundary the initial condition falls in, the decision maker will drift towards one extreme or the other regardless of global efficiency considerations. Under sophisticated long-term planning, the boundary of the basins of attraction is shifted, with the basin of attraction for the efficient optimum becoming larger: for reasonable discount parameters and adjustment costs, agents who start close to the minimum of Π_1 will find the hard path optimal and escape the self-control trap described by the inefficient personal optimum. In the limit as the agent becomes very patient or the costs of adjustment become negligible, even agents who start exactly at the inefficient optimum will eventually converge towards the efficient optimum, as convergence occurs in finite time and the larger payoffs offset the sacrifices made along the way. This, of course, requires unrealistic levels of patience (low discount factor) or self-discipline (low adjustment costs).

5.2 Evolutionary Optimality and Automatic Rationality

On the other extreme of the spectrum, suppose we considered a more biological concept of evolutionary stability. In an implicit evolutionary approach, as in the literature on the evolution of preferences (e.g. Samuelson, 2004; Samuelson and Swinkels, 2006), “nature” acts as a shortcut for the evolutionary process, choosing certain parameters in order to maximize appropriate evolutionary criteria.

Let us start with a simple consideration. A first, *naïve* criterion would be to maximize the probability of an automatic start, z . This objective is often implicitly or explicitly

cited in psychology as a beneficial effect of various forms of interventions. The intuition for why a large proportion of automatic starts is desirable is simple. Since an automatic start does not consume self-control resources, it would be optimal to start automatically as often as possible, hence economizing self-control resources to avert possible breakdowns of will later on.

Since automatic starts are dynamically “trained” by a reinforcement process, intuitively every controlled start, which hence leads to a decision with limited self-control at $t = 2$, will reduce the probability of an automatic start. Thus, the probability of an automatic start will be maximal at the unappealing shirking equilibria. The following result formalizes this intuition.

Proposition 1. *Fix the temptation probability λ . Suppose nature’s objective is to maximize the probability of an automatic start, z .*

- (a) *Suppose that nature is able to select any values of, $d > 0$, $\Delta > 0$, and $g(\Delta) \in]0, 1 - \lambda[$. Then, the optimal values of d , Δ , and $g(\Delta)$ fulfill $\lambda + g(\Delta) > \frac{1}{1+d}$ and hence the unique DPE involves shirking ($x = 0, z = 1 - \lambda$; recall Theorem 1(a)).*
- (b) *Alternatively, suppose $\Delta > 0$, and $g(\Delta) \in]0, 1 - \lambda[$ are also fixed and nature is only able to adjust d . The maximum z is reached at values of d such that $\lambda + g(\Delta) > \frac{1}{1+d}$, and again the unique DPE involves shirking.*

Fixing the temptation parameter λ is the most reasonable approach in this situation. This parameter measures the extent of the basic decision problem we are modeling, which results e.g. due to nature’s tinkering with an immediate consumption bias due to other, more basic survival needs (e.g. the need to store fat for a long winter in a hunter-gatherer society). The remaining parameters have more involved interpretations: Δ is related to the availability of self-control when faced with long-term decisions, $g(\Delta)$ is the reduction in self-control due to a previous exercise thereof, and d is the perceived disutility of failure. Obviously, nature would not be able to set $g(\Delta)$ arbitrarily close to zero due to basic neurophysiological constraints. However, the last result holds independently of such additional constraints as long as $g(\Delta) > 0$.¹⁸

Proposition 1 shows that the mere automatization of behavior cannot have been an evolutionary criterion for the human mind. If we consider Willpower equilibria, Lemma 1 shows that such an evolutionary approach for the automaticity rate Δ is not appropriate, since convexity of Π_1 implies that duality, i.e. $\Delta \in]0, 1[$ cannot be optimal in this sense.

¹⁸Indeed, an all-powerful nature capable of “designing humans” who experience no depletion effects would find that setting $g(\Delta) = 0$ yields $z^*(1) = z^*(0)$, i.e. Willpower equilibria would deliver the same probability of automatic start as Shirking equilibria.

This is, however, not surprising. The reasons for evolutionary optimality of dual processes have to be looked for elsewhere. The main evolutionary advantage of a dual-process decision maker is in *decision speed*, an aspect which is ignored here. Automatic reactions are far quicker than controlled ones. They allowed our ancestors to run first (hence avoiding being eaten) and think later. If you are in the middle of the street and suddenly discover a truck is about to hit you, it is not a good idea to stop and think carefully whether jumping right or left is better. Just jump. Following the evolution parable, nature has probably found it optimal to endow us with a background automaticity rate $\Delta_0 > 0$ to be applied over a variety of different tasks. Once we particularize to a specific task, though, that background automaticity becomes an added constraint and we enter the realm of personal evolution. A related point is made by Samuelson and Swinkels (2006), who point out that visceral urges and impulses reflecting immediate consequences have been shaped by our evolutionary past, and might well create internal conflict with more rational, long-run considerations.

Two comments are in order. First, in spite of these considerations, the “automatization of rationality” remains an attractive goal, and some psychological interventions as implementation intentions (Gollwitzer, 1999) have been designed to pursue it. The interpretation of such interventions, however, is not that they aim to increase the probability of starting automatically, z , for its own sake. In our terms, their aim is, given a DPE (z, x) , to introduce a new automatic process which ignores the outcome of reinforcement and follows the decision that the controlled process would have taken *without* spending self-control resources. Hence, the objective of such interventions is to increase z by changing the nature of the automatic process involved, moving it away from pure reinforcement, and effectively producing more “rational” decisions with a minimum of cognitive costs.

Second, Proposition 1 does not mean that the model at hand or related ones could not be used to tackle questions of evolutionary optimality. In particular, the background rate Δ_0 and the disutility of failure d can only be understood through such an approach.¹⁹ This avenue of research could be pursued by extending our basic game to one of imperfect information where different tasks appear at random, leading e.g. to different values of the temptation parameter. This task is left for future research.

¹⁹Evidence from psychology, incorporated e.g. in prospect theory and reflected in phenomena like loss aversion, shows that negative experiences have a greater impact than positive ones of comparable magnitude (see e.g. Ito, Cacioppo and Lang, 1998; Baumeister et al., 2001). Also, evidence from event-related brain potentials suggests that the brain is wired to react more strongly to negative outcomes than to positive ones (Miltner, Braun and Coles, 1997; Holroyd and Coles, 2002). In other words, even if the objective (in evolutionary terms) utility of failure is equal in magnitude but opposite in sign to the utility of success, nature seems to have found optimal to set the subjective magnitude of failure’s disutility above that of success. This would be an argument for $d > 1$ in the current setting.

6 Conclusion

This paper presents a new, parsimonious modelization of the basic idea that decisions are the result of the interaction of different processes within a decision maker. Drawing on dual-process theories from psychology, which build upon large amounts of empirical evidence but are rarely formalized, we propose to analyze decisions through *Dual-Process Equilibria*, where processes of different nature are modeled in different ways. Controlled processes are modeled as utility-maximizing players from classical decision and game theory. Automatic processes are modeled through techniques and ideas from the literature of learning in games. In the model presented here, and following again insights from psychology, we have identified them with reinforcement processes. The interaction among different processes is captured through an extensive-form formulation analogous to that employed for models with multiple selves.

The first finding is that, if depletion effects are too strong or the disutility of failure is too high, Shirking equilibria appear where decision makers rationally decide not to get started at all. However, psychological interventions might improve welfare by automatizing behavior, i.e. avoiding excessive deliberation.

Willpower equilibria involve controlled processes which start rationally but cause depletion and effortless automatic processes which only start with a certain probability. Inhibiting automatic reactions in favor of controlled ones correspond to acts of willpower. In this setting, the idea of *personal evolution* helps explain individual heterogeneity and shows that suboptimal modes of behavior, as e.g. self-control failures, can correspond to cognitive traps.

We identify two such traps. The first one corresponds to the classical idea of self-control failure, where the individual relies excessively on automatic processes. However, this situation is a local maximum and any attempt to gradually reduce this dependence results in a utility loss, even though a maximum exercise of willpower would be efficient.

The second trap is more surprising. If depletion effects and/or background automaticity are moderate but not too low, the exact opposite might happen. Decision makers might then be caught in a local optimum involving maximum exercise of willpower, which however leads to overall results which are worse than those which could be obtained by relying on automatic (or automatized) behavior only. Hence, decision makers who appear to be doing their best, without however satisfactorily reaching their objectives, might be well advised to radically rethink their approach.

We have identified personal evolution with the gradual adjustment of a self-control parameter, following the direction of increased payoffs. This basic idea could be potentially extended in a number of interesting directions. First, one could consider the possibility

of “jumps” or setbacks in the direction of reduced self-control, or simply postulate that such movements are easier than those in the opposite direction. This would introduce an asymmetry in the concept of personal evolution which might provide an additional explanation why some of the most successful psychological interventions in the realm of self-control operate by automatizing behavior, rather than calling for the exercise of additional willpower. Second, we have modeled the possible breakdown of will down the road through a deterministic automatic process. It would be possible to introduce a probability of breakdown instead, possibly arising from stochastic realizations of temptation values. Associating different realizations to different kinds of problems might help explain why some decision makers face self-control difficulties in certain aspects of their lives and not in others. More generally, the idea of personal evolution proposed here can be applied to more detailed models of self-control problems, incorporating e.g. varying stocks of willpower or tasks with different self-control costs.

References

- Achtziger, Anja, and Carlos Alós-Ferrer.** 2014. “Fast or Rational? A Response-Times Study of Bayesian Updating.” *Management Science*, 60(4): 923–938.
- Achtziger, Anja, and Peter M. Gollwitzer.** 2010. “Motivation and Volition in the Course of Action.” In *Motivation and Action, 3rd Edition.*, ed. Jutta Heckhausen and Heinz Heckhausen, 272–295. UK:Cambridge University Press.
- Achtziger, Anja, Carlos Alós-Ferrer, and Alexander K. Wagner.** 2014. “Money, Depletion, and Prosociality in the Dictator Game.” Working Paper, University of Cologne.
- Achtziger, Anja, Carlos Alós-Ferrer, Sabine Hügelschäfer, and Marco Steinhäuser.** 2014. “The Neural Basis of Belief Updating and Rational Decision Making.” *Social Cognitive and Affective Neuroscience*, 9(1): 55–62.
- Achtziger, Anja, Peter M. Gollwitzer, and Paschal Sheeran.** 2008. “Implementation Intentions and Shielding Goal-Striving From Unwanted Thoughts and Feelings.” *Personality and Social Psychology Bulletin*, 34(3): 381–394.
- Ali, S. Nageeb.** 2011. “Learning Self-Control.” *Quarterly Journal of Economics*, 126: 857–893.
- Alós-Ferrer, Carlos, and Fritz Strack.** 2014. “From Dual Processes to Multiple Selves: Implications for Economic Behavior.” *Journal of Economic Psychology*, 41: 1–11.
- Bargh, John A.** 1989. “Conditional Automaticity: Varieties of Automatic Influences in Social Perception and Cognition.” In *Unintended Thought.*, ed. James S. Uleman and John A. Bargh. Guilford.
- Bargh, John A., and Tanya L. Chartrand.** 1999. “The Unbearable Automaticity of Being.” *American Psychologist*, 54(7): 462–479.
- Baumeister, Roy F.** 2002. “Ego Depletion and Self-Control Failure: An Energy Model of the Self’s Executive Function.” *Self and Identity*, 1(2): 129–136.
- Baumeister, Roy F., Ellen Bratslavsky, C. Finkenauer, and Kathleen D. Vohs.** 2001. “Bad is Stronger than Good.” *Review of General Psychology*, 5: 323–370.
- Baumeister, Roy F., Ellen Bratslavsky, Mark Muraven, and Dianne M. Tice.** 1998. “Ego Depletion: Is the Active Self a Limited Resource?” *Journal of Personality and Social Psychology*, 74(5): 1252–1265.

- Baumeister, Roy F., Kathleen D. Vohs, and Dianne M. Tice.** 2007. “The Strength Model of Self-Control.” *Current Directions in Psychological Science*, 16(6): 351–355.
- Bayer, Ute C., Anja Achtziger, Peter M. Gollwitzer, and Gordon B. Moskowitz.** 2009. “Responding to Subliminal Cues: Do If-Then Plans Cause Action Preparation and Initiation Without Conscious Intent?” *Social Cognition*, 27(2): 183–201.
- Bayer, Ute C., Peter M. Gollwitzer, and Anja Achtziger.** 2010. “Staying on Track: Planned Goal Striving is Protected from Disruptive Internal States.” *Journal of Experimental Social Psychology*, 46(3): 505–514.
- Bechara, Antoine, Xavier Noel, and Eveline A. Crone.** 2006. “Loss of Willpower: Abnormal Neural Mechanisms of Impulse Control and Decision Making in Addiction.” In *Handbook of Implicit Cognition and Addiction.*, ed. Reinout W. Wiers and Alan W. Stacy, 215–232. Thousand Oaks, CA:SAGE.
- Bénabou, Roland, and Jean Tirole.** 2004. “Willpower and Personal Rules.” *Journal of Political Economy*, 112(4): 848–887.
- Benhabib, Jess, and Alberto Bisin.** 2005. “Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions.” *Games and Economic Behavior*, 52(2): 460–492.
- Bernheim, B. Douglas, and Antonio Rangel.** 2004. “Addiction and Cue-Triggered Decision Processes.” *American Economic Review*, 94(5): 1558–1590.
- Börgers, Tilman, and Rajiv Sarin.** 1997. “Learning Through Reinforcement and Replicator Dynamics.” *Journal of Economic Theory*, 77(1): 1–14.
- Botvinick, Mathew M., Jonathan D. Cohen, and Cameron S. Carter.** 2004. “Conflict Monitoring and Anterior Cingulate Cortex: An Update.” *Trends in Cognitive Sciences*, 12: 539–546.
- Bush, R. R., and F. Mosteller.** 1951. “A Mathematical Model for Simple Learning.” *Psychological Review*, 58: 312–323.
- Bush, R. R., and F. Mosteller.** 1955. *Stochastic Models for Learning*. New York:Wiley.
- Camerer, Colin F., and Teck-Hua Ho.** 1999. “Experience-Weighted Attraction Learning in Normal Form Games.” *Econometrica*, 67(4): 827–874.

- Chapman, Simon, and Ross MacKenzie.** 2010. “The Global Research Neglect of Unassisted Smoking Cessation: Causes and Consequences.” *PLoS Medicine*, 7(2): e1000216.
- Charness, Gary, and Dan Levin.** 2005. “When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect.” *American Economic Review*, 95(4): 1300–1309.
- Cross, J. G.** 1973. “A Stochastic Learning Model of Economic Behavior.” *The Quarterly Journal of Economics*, 87: 239–266.
- Cross, J. G.** 1983. *A Theory of Adaptive Economic Behavior*. Cambridge:Cambridge University Press.
- Dane, Erik, and Michael G. Pratt.** 2007. “Exploring Intuition and its Role in Managerial Decision Making.” *Academy of Management Review*, 32(1): 33–54.
- De Neys, Wim, Oshin Vartanian, and Vinod Goel.** 2008. “Smarter than We Think: When Our Brains Detect That We Are Biased.” *Psychological Science*, 19: 483–489.
- Epstein, Seymour.** 1994. “Integration of the Cognitive and the Psychodynamic Unconscious.” *American Psychologist*, 49(8): 709–724.
- Erev, Ido, and Alvin E. Roth.** 1998. “Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria.” *American Economic Review*, 88(5): 848–881.
- Evans, Jonathan St. B. T.** 2008. “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition.” *Annual Review of Psychology*, 59: 255–278.
- Ferguson, Stuart G., Saul Shiffman, Joseph G. Gitchell, Mark A. Sembower, and Robert West.** 2009. “Unplanned Quit Attempts—Results from a U.S. Sample of Smokers and Ex-Smokers.” *Nicotine and Tobacco Research*, 11(7): 827–832.
- Frank, Michael J., Michael X. Cohen, and Alan G. Sanfey.** 2009. “Multiple Systems in Decision Making: A Neurocomputational Perspective.” *Current Directions in Psychological Science*, 18(2): 73–77.
- Fudenberg, Drew, and David K. Levine.** 2006. “A Dual-Self Model of Impulse Control.” *American Economic Review*, 96(5): 1449–1476.

- Fudenberg, Drew, and David K. Levine.** 2012. "Timing and Self-Control." *Econometrica*, 80(1): 1–42.
- Fujita, Kentaro.** 2011. "On Conceptualizing Self-Control as More Than the Effortful Inhibition of Impulses." *Personality and Social Psychology Review*, 15(4): 352–366.
- Fujita, Kentaro, and H. Anna Han.** 2009. "Moving Beyond Deliberative Control of Impulses: The Effect of Construal Levels on Evaluative Associations in Self-Control Conflicts." *Psychological Science*, 20(7): 799–804.
- Gailliot, Matthew T., Roy F. Baumeister, C. Nathan DeWall, Jon K. Maner, E. Ashby Plant, Dianne M. Tice, Lauren E. Brewer, and Brandon J. Schmeichel.** 2007. "Self-Control Relies on Glucose as a Limited Energy Source: Willpower Is More Than a Metaphor." *Journal of Personality and Social Psychology*, 92(2): 325–336.
- Gollwitzer, Peter M.** 1999. "Implementation Intentions: Strong Effects of Simple Plans." *American Psychologist*, 54: 493–503.
- Gollwitzer, Peter M., and Paschal Sheeran.** 2006. "Implementation Intentions and Goal Achievement: A Metaanalysis of Effects and Processes." *Advances in Experimental Social Psychology*, 38: 69–119.
- Hagger, Martin S., Chantelle Wood, Chris Stiff, and Nikos L. D. Chatzisarantis.** 2010. "Ego Depletion and the Strength Model of Self-Control: A Meta-Analysis." *Psychological Bulletin*, 136(4): 495–525.
- Holroyd, Clay B., and Michael G. H. Coles.** 2002. "The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity." *Psychological Review*, 109(4): 679–709.
- Inzlicht, Michael, and Brandon J. Schmeichel.** 2012. "What Is Ego Depletion? Toward a Mechanistic Revision of the Resource Model of Self-Control." *Perspectives on Psychological Science*, 5(7): 450–463.
- Ito, Tiffany A., John T. Cacioppo, and Peter J. Lang.** 1998. "Eliciting Affect Using the International Affective Picture System: Bivariate Evaluation and Ambivalence." *Personality and Social Psychology Bulletin*, 24: 855–879.
- Kahneman, Daniel.** 2003. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *American Economic Review*, 93(5): 1449–1475.

- Kahneman, Daniel.** 2011. *Thinking, Fast and Slow*. New York:Farrar, Straus and Giroux.
- Lieberman, Matthew D.** 2003. “Reflexive and Reflective Judgment Processes: A Social Cognitive Neuroscience Approach.” In *Social Judgments: Implicit and Explicit Processes*. , ed. Joseph P. Forgas, Kipling D. Williams and W. von Hippel, 44–67. UK:Cambridge University Press.
- Lieberman, Matthew D., Johanna M. Jarcho, and Ajay B. Satpute.** 2004. “Evidence-Based and Intuition-Based Self-Knowledge: An fMRI Study.” *Journal of Personality and Social Psychology*, 87(4): 421–435.
- Loewenstein, George, and Ted O’Donoghue.** 2005. “Animal Spirits: Affective and Deliberative Processes in Economic Behavior.” Working Paper, Carnegie Mellon University.
- Masicampo, Emer J., and Roy F. Baumeister.** 2008. “Toward a physiology of dual-process reasoning and judgment: Lemonade, willpower, and effortful rule-based analysis.” *Psychological Science*, 19(3): 255–260.
- McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen.** 2004. “Separate Neural Systems Value Immediate and Delayed Monetary Rewards.” *Science*, 306(5695): 503–507.
- Mead, Nicole L., Roy F. Baumeister, Francesca Gino, Maurice E. Schweitzer, and Dan Ariely.** 2009. “Too Tired to Tell the Truth: Self-Control Resource Depletion and Dishonesty.” *Journal of Experimental Social Psychology*, 45(3): 594–597.
- Milkman, Katherine L., John Beshears, James J. Choi, David Laibson, and Brigitte C. Madrian.** 2011. “Using implementation intentions prompts to enhance influenza vaccination rates.” *Proceedings of the National Academy of Sciences of the USA*, 108(26): 10415–10420.
- Miltner, Wolfgang H. R., Christoph H. Braun, and Michael G. H. Coles.** 1997. “Event-Related Brain Potentials Following Incorrect Feedback in a Time-Estimation Task: Evidence for a “Generic” Neural System for Error Detection.” *Journal of Cognitive Neuroscience*, 9: 788–798.
- Molden, Daniel C., Chin Ming Hui, Abigail A. Scholer, Brian P. Meier, Eric E. Noreen, Paul R. D’Agostino, and Valerie Martin.** 2012. “Motivational Versus

- Metabolic Effects of Carbohydrates on Self-Control.” *Psychological Science*, 23: 1137–1144.
- Moore, Bert, Walter Mischel, and Antonette Zeiss.** 1976. “Comparative Effects of the Reward Stimulus and its Cognitive Representation in Voluntary Delay.” *Journal of Personality and Social Psychology*, 34(3): 419–424.
- Muraven, Mark, Dianne M. Tice, and Roy F. Baumeister.** 1998. “Self-Control as Limited Resource: Regulatory Depletion Patterns.” *Journal of Personality and Social Psychology*, 74(3): 774–789.
- Nickerson, David W., and Todd Rogers.** 2010. “Do You Have a Voting Plan?: Implementation Intentions, Voter Turnout, and Organic Plan Making.” *Psychological Science*, 21(2): 194–199.
- O’Donoghue, Ted, and Matthew Rabin.** 1999. “Doing It Now or Later.” *American Economic Review*, 89(1): 103–124.
- O’Donoghue, Ted, and Matthew Rabin.** 2001. “Choice and Procrastination.” *Quarterly Journal of Economics*, 116(1): 121–160.
- Ozdenoren, Emre, Stephen Salant, and Dan Silverman.** 2012. “Willpower and the Optimal Control of Visceral Urges.” *Journal of the European Economic Association*, 10(2): 342–368.
- Ridderinkhof, K. Richard, Markus Ullsperger, Eveline A. Crone, and Sander Nieuwenhuis.** 2004. “The Role of the Medial Frontal Cortex in Cognitive Control.” *Science*, 306: 443–447.
- Samuelson, Larry.** 2004. “Information-Based Relative Consumption Effects.” *Econometrica*, 72(1): 93–118.
- Samuelson, Larry, and Jeroen Swinkels.** 2006. “Information, Evolution and Utility.” *Theoretical Economics*, 1: 119–142.
- Sanfey, Alan G., James K. Rilling, Jessica A. Aronson, Leigh E. Nystrom, and Jonathan D. Cohen.** 2003. “The Neural Basis of Economic Decision-Making in the Ultimatum Game.” *Science*, 300(5626): 1755–1758.
- Schelling, Thomas C.** 1984. “Self-Command in Practice, in Policy, and in a Theory of Rational Choice.” *American Economic Review*, 74(2): 1–11.

- Siegel, S., and D. A. Goldstein.** 1959. "Decision Making Behavior in a Two-Choice Uncertain Outcome Situation." *Journal of Experimental Psychology*, 57(1): 37–42.
- Strack, Fritz, and Roland Deutsch.** 2004. "Reflective and Impulsive Determinants of Social Behavior." *Personality and Social Psychology Review*, 8(3): 220–247.
- Thaler, Richard H., and Hersch M. Shefrin.** 1981. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89(2): 392–406.
- Vohs, Kathleen D., and Ronald J. Faber.** 2007. "Spent Resources: Self-Regulatory Resource Availability Affects Impulse Buying." *Journal of Consumer Research*, 33: 537–547.
- Vohs, Kathleen D., and Todd F. Heatherton.** 2000. "Self-regulatory failure: A resource-depletion approach." *Psychological Science*, 11(3): 249–254.
- Weber, Elke U., and Eric J. Johnson.** 2009. "Mindful Judgment and Decision Making." *Annual Review of Psychology*, 60: 53–85.
- West, Robert, and Taj Sohal.** 2006. "'Catastrophic' Pathways to Smoking Cessation: Findings from National Survey." *British Medical Journal*, 25(332): 458–460.

APPENDIX: Proofs

Proof of Lemma 1. For notational convenience, let $R_\Delta = (1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g$. R_Δ is the discriminant of the polynomial we analyzed to compute z^* (recall (6)), for $x = 1$. As we argued in the main text, $R_\Delta > 0$.

To see that Π_1 is a strictly convex function of Δ , we compute

$$(A1) \quad \frac{\partial \Pi_1}{\partial \Delta} = (1 - \lambda(1 + d)) \left[-\frac{1}{2} \left[\lambda + \frac{(1 - \Delta\lambda)\lambda + 2(1 - 2\Delta)g}{\sqrt{R_\Delta}} \right] + \left(\frac{g(1 + d)}{(1 - \lambda(1 + d))} \right) \right]$$

$$\frac{\partial^2 \Pi_1}{\partial \Delta^2} = -\frac{1}{2} \frac{(1 - \lambda(1 + d))}{R_\Delta} \left[(-\lambda^2 - 4g) \sqrt{R_\Delta} + \frac{((1 - \Delta\lambda)\lambda + 2(1 - 2\Delta)g)^2}{\sqrt{R_\Delta}} \right]$$

Recall that, by (2), $\lambda(1 + d) < 1$. Then,

$$\begin{aligned} \frac{\partial^2 \Pi_1}{\partial \Delta^2} \underset{\leq}{\geq} 0 &\iff (-\lambda^2 - 4g) \sqrt{R_\Delta} + \frac{((1 - \Delta\lambda)\lambda + 2(1 - 2\Delta)g)^2}{\sqrt{R_\Delta}} \underset{\leq}{\geq} 0 \iff \\ &((1 - \Delta\lambda)\lambda + 2(1 - 2\Delta)g)^2 \underset{\leq}{\geq} (\lambda^2 + 4g) R_\Delta \iff \\ (1 - 2\Delta)^2 g^2 + (1 - \Delta\lambda)\lambda(1 - 2\Delta)g &\underset{\leq}{\geq} -\lambda^2 \Delta(1 - \Delta)g + g(1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g^2 \\ &\iff g \underset{\leq}{\geq} 1 - \lambda \end{aligned}$$

Since $g < 1 - \lambda$, we conclude that $\frac{\partial^2 \Pi_1}{\partial \Delta^2} > 0$ for all Δ , and the claim holds. \square

Proof of Lemma 2: First note that if $\frac{\partial \Pi_1}{\partial \Delta} \Big|_{\Delta=0} \geq 0$, then Π_1 is strictly increasing in Δ by Lemma 1 and reaches a maximum at $\Delta = 1$. This happens if and only if (evaluating (A1) at $\Delta = 0$)

$$\frac{\partial \Pi_1}{\partial \Delta} \Big|_{\Delta=0} \geq 0 \iff \lambda + g \leq \left(\frac{g(1 + d)}{(1 - \lambda(1 + d))} \right) \iff g \geq \lambda \frac{1 - \lambda(1 + d)}{d + \lambda(1 + d)} =: \bar{g}$$

Note that $\bar{g} \in]0, 1[$.

If $\frac{\partial \Pi_1}{\partial \Delta} \Big|_{\Delta=1} \leq 0$, then Π_1 is strictly decreasing in Δ and reaches a maximum at $\Delta = 0$.

This happens if and only if (evaluating (A1) at $\Delta = 1$)

$$\begin{aligned} \frac{\partial \Pi_1}{\partial \Delta} \Big|_{\Delta=1} \leq 0 &\iff \lambda + \frac{(1-\lambda)\lambda - 2g}{(1-\lambda)} \geq \left(\frac{2g(1+d)}{(1-\lambda)(1+d)} \right) \iff \\ &(1-\lambda(1+d))(\lambda(1-\lambda) - g) \geq g(1-\lambda)(1+d) \iff \\ (1-\lambda(1+d))\lambda(1-\lambda) &\geq g[(1-\lambda)(1+d) + (1-\lambda(1+d))] = g[1 + (1-2\lambda)(1+d)] \\ &\iff g \leq \lambda(1-\lambda) \frac{1-\lambda(1+d)}{1+(1-2\lambda)(1+d)} =: \underline{g} \end{aligned}$$

Notice that $\lambda < \frac{1}{1+d}$ implies that the denominator in the last fraction is always strictly positive.

It can be shown that $\underline{g} < 1 - \lambda$. This holds if and only if

$$\begin{aligned} \lambda(1-\lambda(1+d)) < 1 + (1-2\lambda)(1+d) &\iff \lambda < 1 + (1-\lambda)^2(1+d) \\ &\iff 1 + (1-\lambda)(1+d) > 0 \end{aligned}$$

which proves the claim. Notice that $\underline{g} < \bar{g}$ follows by convexity of Π_1 (Lemma 1), but it also can be easily checked directly. If $\underline{g} < g < \bar{g}$, then there exists an interior *minimum*, but still there is a corner solution. \square

Proof of Theorem 2. The proof relies on Lemmata 1 and 2. First, we claim that there exists g^* such that $\Pi_1|_{\Delta=0} \leq \Pi_1|_{\Delta=1} \iff g \geq g^*$.

To see this, note that since Π_1 is strictly convex by Lemma 1, the function Π_1 has either a global maximum at $\Delta = 0$, or a global maximum at $\Delta = 1$. Whether the first or the second case occur depends only on the following inequality (recall (12)):

$$\Pi_1|_{\Delta=0} \leq \Pi_1|_{\Delta=1} \iff \frac{g(1+d)}{(1-\lambda(1+d))} \geq \lambda \iff g \geq \lambda \left[\frac{1}{1+d} - \lambda \right] =: g^*$$

which proves the claim. Note that if $g = g^*$ (hence there are two global maxima), by Lemma 2 we must be in the case $\underline{g} < g < \bar{g}$, hence $\underline{g} < g^* < \bar{g}$

We now turn to a case-by-case analysis, systematically following the possible values of the parameters.

Case 1. If $g \leq \underline{g}$, then utility would be maximized at full control, $\Delta = 0$; further, utility is strictly decreasing in Δ by Lemma 2. Hence maximum control, $\Delta = \Delta_0$, is the only personal optimum, which is also efficient. This is Case (I) in the statement.

Case 2. If $g \geq \bar{g}$, then utility is maximized at full automaticity, $\Delta = 1$; further, utility is strictly increasing in Δ by Lemma 2. Hence full automaticity is the only personal

optimum, which is also efficient. This is case (IIa) in the statement.

Case 3. If $g^* \leq g < \bar{g}$, then there is an interior *minimum* at $\Delta = \Delta(g)$ by Lemma 2 and $\Delta = 1$ is the global maximum by the claim above.²⁰ Then, we have to distinguish two subcases.

Case 3a. If $\Delta_0 \geq \Delta(g)$, then utility is strictly increasing at Δ_0 and again full automaticity is efficient and the only personal optimum.²¹ This belongs to Case (IIb).

Case 3b. If $\Delta_0 < \Delta(g)$, then utility is strictly decreasing at Δ_0 . It follows that maximum control, $\Delta = \Delta_0$ is a personal optimum but full automaticity is efficient. This belongs to Case (IV).

Case 4. If $\underline{g} < g < g^*$, then there is an interior *minimum* at $\Delta = \Delta(g)$ by Lemma 2 and $\Delta = 0$ is the global maximum by the claim above. It follows that there exists $\Delta'(g) \in]0, \Delta(g)[$ such that

$$\Pi_1|_{\Delta=\Delta_0} \begin{matrix} \geq \\ \leq \end{matrix} \Pi_1|_{\Delta=1} \iff \Delta_0 \begin{matrix} \leq \\ \geq \end{matrix} \Delta'(g)$$

Then, we have to distinguish several subcases.

Case 4a. If $\Delta_0 \geq \Delta(g)$, then utility is strictly increasing at Δ_0 and again full automaticity is efficient and the only personal optimum. This belongs again to Case (IIb).

Case 4b. If $\Delta'(g) < \Delta_0 < \Delta(g)$, then utility is strictly decreasing at Δ_0 ; further, the utility at Δ_0 is strictly smaller than that at $\Delta = 1$. Hence, full automaticity is the efficient optimum, but maximum control is an inefficient personal optimum. This belongs again to Case (IV).

Case 4c. If $\Delta_0 = \Delta'(g)$, then utility is strictly decreasing at Δ_0 ; further, the utility at Δ_0 is equal to that at $\Delta = 1$. This is the knife-edge Case (V).

Case 4d. If $\Delta_0 < \Delta'(g)$, then utility is strictly decreasing at Δ_0 , and the utility at Δ_0 is strictly larger than that at $\Delta = 1$. Hence, maximum control is the efficient optimum, but full automaticity is an inefficient personal optimum. This is Case (III).

Continuity of $\Delta(\cdot)$ and the facts that $\Delta(\underline{g}) = 1$ and $\Delta(\bar{g}) = 0$ follow from the definition of $\Delta(\cdot)$ and Lemma 2. Continuity of $\Delta'(g)$ and the facts that $\Delta'(g) = 1, 0 \leq \Delta'(g) < \Delta(g)$ on $g \in]\underline{g}, \bar{g}[$, and $\Delta'(g^*) = 0$ follow from definition of $\Delta'(\cdot)$. \square

Proof of Proposition 1. (a) Notice that d does not appear in any equilibrium computation. It merely determines which type of equilibrium does the model have, Hence, selecting d amounts to selecting the equilibrium type. Given an equilibrium, the relationship between the probability of a controlled start, x , and that of an automatic start, z , is given

²⁰If $g = g^*$, there are two global maxima at $\Delta = 0, 1$, but the first is irrelevant since $\Delta_0 > 0$.

²¹In the corner case $\Delta_0 = \Delta(g)$, utility is also strictly increasing given the constraint $\Delta \geq \Delta_0$.

by equation (6). A simple computation shows that

$$(A2) \quad 2\Delta \frac{dz^*(x)}{dx} = (1 - \Delta) \left[-1 + \frac{(1 - \Delta)x + \Delta(1 - \lambda) - 2\Delta g(\Delta)}{\sqrt{((1 - \Delta)x + \Delta(1 - \lambda))^2 - 4x\Delta(1 - \Delta)g(\Delta)}} \right]$$

Consider the fraction appearing in (A2). The square of this fraction is strictly smaller than one (and hence so is its absolute value) if and only if

$$\begin{aligned} (((1 - \Delta)x + \Delta(1 - \lambda)) - 2\Delta g(\Delta))^2 &< ((1 - \Delta)x + \Delta(1 - \lambda))^2 - 4x\Delta(1 - \Delta)g(\Delta) \\ \iff 4\Delta^2 g(\Delta)^2 - 4((1 - \Delta)x + \Delta(1 - \lambda))\Delta g(\Delta) &< -4x\Delta(1 - \Delta)g(\Delta) \\ &\iff g(\Delta) < 1 - \lambda \end{aligned}$$

and this last condition is always fulfilled (recall that $1 - \lambda - g(\Delta) > 0$ is the probability of the controlled process determining the decision at $t = 2$). Hence, it follows from (A2) that $z^*(x)$ is strictly decreasing and the conclusion follows.

(b) For each combination of the parameters Δ , $g(\Delta)$, and d , by Theorem 1 three situations might result. If $\lambda + g(\Delta) > \frac{1}{1+d}$, there is a unique DPE of the Shirking type with $z = 1 - \lambda$. If $\lambda + g(\Delta) < \frac{1}{1+d}$, there is a unique DPE of the Willpower type with z given by (9). If $\lambda + g(\Delta) = \frac{1}{1+d}$, there is a continuum of equilibria with $x \in [0, 1]$ and $z(x)$ given by (6).

Consider a situation in the second case and let $z^*(1)$ denote the associated probability of an automatic start. Increasing d leaving other parameters constant (for example) leads to an outcome in the first case, with a probability of automatic start $z^*(0) = 1 - \lambda$. But then, using (9),

$$\begin{aligned} z^*(0) &> z^*(1) \iff \\ 1 - \lambda &> 1 + \frac{1}{2\Delta} \left[-(1 + \Delta\lambda) + \sqrt{(1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g(\Delta)} \right] \iff \\ 1 - \Delta\lambda &> \sqrt{(1 - \Delta\lambda)^2 - 4\Delta(1 - \Delta)g(\Delta)} \iff \\ 0 &< 4\Delta(1 - \Delta)g(\Delta) \end{aligned}$$

which always holds independently of Δ , $g(\Delta)$, and d . Hence, all Willpower equilibria have smaller z than Shirking equilibria.

Suppose now that nature sets Δ , $g(\Delta)$, and d such that there is a continuum of mixed-strategy equilibria. The computations in part (a) show that the equilibrium which maximizes z is the one with complete shirking, $x = 0$, and the probability of automatic start is identical to that of shirking equilibria. This completes the proof. \square