How Our Commitments Slip Away From Us

Alexey Guzey*

June 20, 2018
Click here for the latest version.

Abstract

Much economic research is concerned with the question of why our commitments to ourselves fail, whether it's lack of self-confidence, lack of willpower in the face of temptations, conflict between the "selves", or one of many other explanations. We introduce and formalize a new model of commitment failure. We posit that the subjective experience of failing a commitment is frequently indistinguishable from the subjective experience of following it. Thus, we model a person consisting of two selves: Planner (she) and Doer (he). Planner devises the rule for the person's long-term benefit and Doer follows his perception of Planner's rule. We analyze the conditions under which Planner is (not) able to devise an effective rule. We show that under some conditions, Doer, while believing that he continues to follow the Planner's rule, descends into complete procrastination, while, under some conditions, Doer always works to the person's long-term benefit.

*I am grateful to Alexis Belianin for his patient, yet stimulating guidance and to Anastasia Kuptsova for many prolific discussions and continued support.

Contents

1 Introduction

People plan ahead. Frequently, they are aware that their future selves may not be so keen on fulfilling some of their past selves' desires. Thus, people are searching for ways to circumvent their future selves' (1) discretionary ability and (2) to alter their future selves' incentives (see [1], [2] for an overview of self-regulation literature). Economists first started to pay attention to the problem of temporal inconsistency with Strotz [3], followed by seminal works by Thaler [4] and Loewenstein [5].

Many approaches to the problem of self-regulation have since then appeared in the economic literature (see [6], [7] for an overview). Some of the most influential are: hyperbolic and quasi-hyperbolic discounting ([8], [9]), temptation ([10]–[12]), Multiple selves ([4], [13]–[15]), self-signaling [16].

Most prominently, we lean on the idea of "personal rules" ([17]–[19]), which act as internal commitment mechanisms, modifying future behavior. These commitments are not perfect, and in contrast to external commitment mechanisms, can be broken, which has lead many researchers to investigate the conditions under which we break personal rules.

However, instead of trying to explain why we *break* personal rules, we posit that the subjective experience of personal rule breaking is almost never of breaking rules (something Loewenstein [14] alluded to while talking about addiction). As Ainslie [17] acutely notices, the rules are broken by us progressively acquiescing and accepting larger and larger excuses for why the rule is not applicable in this particular situation (p. 87):

> [T]he same principle that keeps Thanksgiving from setting a precedent might also work for my birthday and, with decreasing credibility, for the Fourth of July, St. Patrick's Day, Labor Day, Arbor Day, St. Swithin's Day, and Just This Once. This kind of logic can degrade a personal rule without my ever breaking it.

The subjective experience of having broken a commitment is us suddenly finding ourselves no longer following it, having no idea as to how exactly it happened.

Following the literature ([4], [15]), we model the person as consisting of two selves: Planner (she) and Doer (he). However, in contrast with the majority of the literature, we do not assume a conflict between the two selves. Doer *is*, in fact, trying to follow Planner's rules. He is just not very good at it.

For example, Doer feels that procrastinating a little bit and then starting to work is very similar to just working, as Planner might have prescribed. Thus, Doer does procrastinate a little bit but still feels that he followed Planner's rule. Due to imperfect memory, in the next period, Doer can only remember that (1) he procrastinated a little bit and that (2) he followed Planner's rule, which morphs Doer's perception of the rule, towards a little bit of procrastination being allowed.

Over time, Doer may start feeling that procrastinating more and more, and working less and less, while still continuing to feel that he follows the rule.

Depending on Planner's sophistication, which we model as access to Doer's mental processes, in some cases, she may devise a rule, which would prevent this dissolution from happening. For example, Planner may strategically lure Doer to procrastinate in a way that eventually leads him to return to work, rather than procrastinate *ad infinitum*.

To the best of our knowledge, no theoretical model of slipping away from commitments while maintaining the belief in rule following has been provided yet. Thus, we attempt to create a model which would explain which conditions lead to this phenomenon and to provide directions on what we can do to prevent it from happening.

2 Literature Review

The model I build is most closely related to work by Mazar et al. [20]. They suggest that offered with an opportunity to cheat, people will balance their desire to seem honest to themselves with the financial gains from cheating, which will result in people cheating, but only as long as they are able to retain their positive self-concept, which is similar to our claim that people will deviate from rules, but only as long as they feel that they are still following them. However, Mazar et al. do not build a formal model of their idea. Moreover, one of their key claims is that the extent of cheating people will engage in is naturally bounded by the point at which they will no longer be able to maintain a positive self-concept. In contrast, I show that under some conditions, the person's behavior will become indistinguishable from them acting, following the analogy, "completely dishonest".

A different strand of psychological research related to my model is the study of the "foot-in-the-door" technique, introduced by Freedman [21]. The technique suggests that in order to make somebody acquiesce to a large demand one should first start with a small demand (which will surely be accepted) and then gradually increase the size of the demand, until the person is ready to acquiesce to the demand they would have rejected initially (for an overview of research on this technique see [22]). I propose that the brain is essentially using the "foot-in-the-door" technique on its owner, by way of adjusting the perception of rules set by them little by little, resulting in the possibility of large differential between the perception of rules and actual rules, with us not noticing it.

Economists first started to pay attention to time-inconsistency of people's preferences with work by Strotz 1955 who noted that intertemporal conflict between selves arises if the person's discounting rate is not exponential, and suggested that the solution is to either impose external restrictions on behavior or to only make plans which are time-consistent.

Thaler and Shefrin [4] were first to formalize the idea of multiple (in their case two) selves. Two selves are a planner (who is interested in the long-term maximization of utility) and a doer (who is only interested in current consumption), who are analogue to a principal and an agent, in a principal-agent model. The tools the planner uses are incentives altering and rule setting.

Many alternative interpretations of multiple selves have followed [13]–[15], some of them continuing the principal-agent analogy [23].

Economic literature most closely related to our paper examines ways people can modify their future behaviors only using internal commitments. Tirole and Bénabou [18] built one of the most influential such models. They argue that the fear for one's self-reputation over one's willpower is the key mechanism for the effectiveness of promises to oneself, or as authors call them, "personal rules".

Finally, the similarity literature explores many of the same themes as we do

[24], [25]. In particular, Gilboa and Schmeidler's [26] case-based decision theory, which contends that the decisions are made based on a recollection of similar past cases and outcomes that followed the decisions in them. We share their opinion that the past decisions influence us directly, rather than through more cognitive mechanisms, as most other researchers assert.

3 Twitter and Facebook Example

The following example is meant to informally introduce the key ideas of the paper. We will refer to it frequently in the future.

Suppose there are 5 periods, each 25-minute long. The person would prefer to work (w) for the entire duration of each period.

However, they could also spend 5 (10, 15) minutes on Facebook and then work 20 (15, 10) minutes until the end of the period (f_5 (f_{10}, f_{15})) (they cannot spend more than 20 minutes on Facebook, because they have few friends and the News Feed runs out).

Alternatively, the person could also spend 5 (10, 15, 20 or 25) minutes on Twitter and then work 20 (15, 10, 5, 0) minutes until the end of the period (t_5 (t_{10}, t_{15}, t_{20}, t_{25})) (since they follow a lot of people and the feed is effectively infinite),

The person dislikes work and enjoys browsing Facebook, and enjoys browsing Twitter even more. Thus, short-term utility U is defined as follows:
$U(w) = 0$
$U(f_5) = 1, U(f_{10}) = 3, U(f_{15}) = 5$
$U(t_5) = 2,\ U(t_{10}) = 4, U(t_{15}) = 6,\ U(t_{20}) = 8,\ U(t_{25}) = 10$

Now, the suppose the person devises the following rule for themselves: *I'm only allowed to work.* However, when the first period starts and the person has to decide what to do, they do not act solely upon their rule. In the moment, they act upon their *perception* of the rule. In the moment, it may seem that the action "work for 25 minutes" is very similar to "spend 5 minutes on Facebook and then work 20 minutes" and to "spend 5 minutes on Twitter and then work 20 minutes". Since $U(t_5)$ is highest among the options, the person decides to pursue t_5. It is important to note that the person *does not feel* as if they've broken the rule.

In the next period, the person's perception of the rule is adjusted. Due to imperfect recall, they only remember that they (1) felt that they were following the rule and (2) that they spent 5 minutes on Twitter, which is now perceived as having been allowed. So, now they believe that w is allowed and t_5 is allowed. Repeating the thought process from above, the agent now decides to pursue t_{10} (since it feels very similar to t_5, which is "allowed", and it brings the highest utility from all allowed behaviors).

It is easy to see that by the last period the person will end up browsing Twitter the entire time and not working at all. At this point, they might finally realize that the rule was broken, but they have little idea of how this happened since they never felt that they were breaking it.

4 The Model

We consider a discrete-time model with $n \in \mathbb{N}$ periods and a finite set of possible actions X, consisting of action w (work for the entire period), and, possibly, other actions (which may partially include work, like e.g. t_5 in example above).

For each $x_i \in X$, define the set S_i, consisting of actions similar to x_i, with the only restriction being $x_i \in S_i \ \forall i$.

Let S be the collection of all S_i.

We model the person as consisting of two selves: Planner (she) and Doer (he).

Planner devises the rule R, defined as a non-empty set of allowed actions x_i, maximizing the amount of work she expects Doer to perform following the rule. Planner does not possess the ability to notice when Doer does not follow her rule until after the last period.

Doer maximizes his utility U, while following his perception of the rule. He perceives that the rule is followed if the action he takes are similar to actions he believes to be allowed. Doer possesses imperfect recall, which means he infers the rule from his actions last period.

We assume that Doer has perfect knowledge of S.

Let the utility Doer gains from any action to be from [0,1], with work utility $U(w) = 0$, procrastination utility $U(p) = 100$ and $0 < U(x) < 100, \ \forall x \neq w, \ p$, meaning that Doer would prefer doing anything else to work, and would prefer procrastinating to anything else. Further, $U(x) = U(y) \implies x = y$, i.e. no two actions yield the same utility.

We will now see how Planner's sophistication affects the efficacy of the rule she can devise to make Doer act in accordance with the person's long-term goals.

4.1 Naive Planner

Naive Planner does not observe Doer's mental processes, which means she does not realize the existence of S and believes that Doer follows the rule strictly.

Thus, Naive Planner devises the following rule, trying to maximize the amount of time Doer works:

Naive Rule $R_{\text{Naive}} = \{w\}$

Proof. Suppose $x_i \neq w, \ x_i \in R$. Then, since $U(x_i) > U(w) \ \forall x_i \neq w$, will lead to Doer choosing x_i instead of w at least once, which will decrease the amount of time Doer works. Thus, $R_{\text{Naive}} = \{w\}$ is optimal. □

Consequently, according to Naive Planner's model of Doer, he will always be working.

Reminder: w is work, S_i is the set of actions similar to x_i, the rule R is a non-empty set of allowed actions x_i

Doer's actual behavior depends on the S. First, suppose, work is dissimilar from any other behavior, i.e. $S_w = w$.

Reminder: $U(w) = 0$, $U(x) > 0$ $\forall x \neq w$

Then, work is the only available action for Doer and he will only work, just as Planner expected.

However, suppose work is similar to some other actions. We would like to know what Doer will do on this case in each round. To find out, we construct a directed graph $G(I, g)$, consisting of a set of vertices $I = \{v_1, ..., v_k\}$, where k is the number of actions on X, so that each vertex v_i corresponds to an action x_i, and a $n \times n$ adjacency matrix g, where $g_{ij} = 1$ if there exists an edge from v_i to v_j and 0 otherwise.

To create edges we will need to find the best similar action for every action x_i, which we will denote b_i. Action $x_j \in S_i$ is the best similar action for action x_i if:

Reminder: S_i is the set of actions similar to x_i

$$U(x_j) = \max_{x \in S_i}\{U(x)\}$$

Then, for each vertex v_i create an edge to the vertex, corresponding to the best similar to x_i action.

Now, Doer starts at w and with each round moves along the path, if one exists. This is equivalent to Doer choosing the action with the highest utility among behaviors he believes to be allowed. Actions corresponding to already visited nodes are perceived as allowed.

Henceforth, we will base our discussion of Doer's behavior based on a G, generated as described above, since G fully describes Doer's behavior.

For an example of generated graph, see figure below.

Proposition 4.1. *G is acyclic.*

Proof. By construction, there exists a path from v_i to v_j only if $U(x_i) < U(x_j)$. Thus, if there exists a cycle which includes v_i, then $U(x_i) > U(x_i)$. Contradiction. □

Corollary 4.1.1. *Provided a sufficiently large number of periods, Doer will reach the action from which he will not deviate.*

Reminder: p is complete procrastination

Note, this does not imply that the sole existence of a path from w to p will lead to Doer eventually moving to p.

Thus, in contrast with Naive Planner's prediction, Doer is not guaranteed to follow the rule.

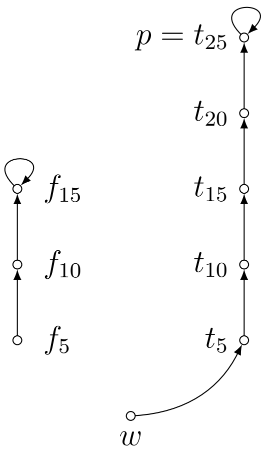

fig. 1

Example. Recall an informal example with Twitter and Facebook (we will refer to it as TF). G_{TF} is shown in fig. 1.
$U(w) = 0$
$U(f_5) = 1, U(f_{10}) = 3, U(f_{15}) = 5$
$U(t_5) = 2,\ U(t_{10}) = 4, U(t_{15}) = 6,\ U(t_{20}) = 8,\ U(t_{25}) = 10$
$S_w = \{w,\ f_5,\ t_5\}$
$S_{f_5} = \{w,\ f_5\ f_{10}\},\ S_{f_{10}} = \{f_5,\ f_{10}, f_{15}\},\ S_{f_{15}} = \{f_{10},\ f_{15}\}$
$S_{t_5} = \{w,\ t_5\ t_{10}\},\ S_{t_{10}} = \{t_5,\ t_{10}, t_{15}\},\ S_{t_{15}} = \{t_{10},\ t_{15}, t_{20}\},\ S_{t_{20}} = \{t_{15},\ t_{20},\ t_{25}\},\ S_{t_{25}} = \{t_{20},\ t_{25}\}$

Enacting R_{Naive}, in the first period, Doer would choose action t_5; in the second: t_{10}; in the third: t_{15}; in the forth: t_{20}; in the fifth: t_{25}. Thus, total work is 50 minutes. On the graph:

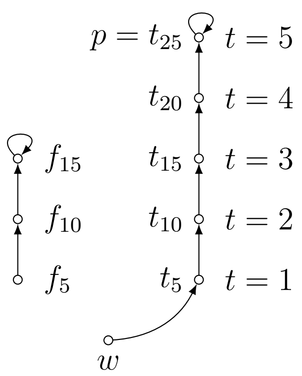

fig. 2

4.2 Sophisticated Planner

Sophisticated Planner has access to Doer's mental processes, i.e. she both knows S and that Doer's behavior is completely determined by S (represented by G). Thus, she can try to devise the rule to take advantage of this knowledge.

Reminder: S_i is the set of actions similar to x_i, S is the collection of all S_i

Let's return to the TF example. Here, instead of forbidding everything except work, Planner could introduce the following rule, allowing browsing Facebook for 5 minutes:

Sophisticated Rule $R_{\text{Sophisticated}} = \{f_5\}$

In the first period, Doer would choose action f_{10}; in the second: f_{15}; in the third: f_{15}; in the forth: f_{15}; in the fifth: f_{15}. Thus, total work is 55 minutes, compared to R_{Naive}'s 50 minutes. On the graph:

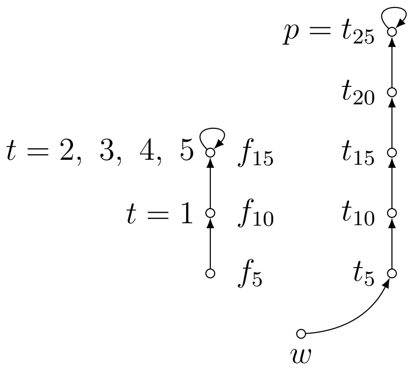

fig. 3

Proposition 4.2. *$R_{Sophisticated} = \{f_5\}$ is optimal in TF.*

Proof. Note, that any rule, which results in Doer choosing among t actions, brings less or equal amount work than R_{Naive} (examine fig. 2). Further, note, that any rule, which results in Doer choosing among f actions, brings less or equal amount of work than $R_{\text{Sophisticated}}$ (examine fig. 3). Thus, $R_{\text{Sophisticated}}$ is optimal. $\square$

Proposition 4.3. *Any rule consisting of more than one action has an equivalent rule consisting of exactly one action.*

Proof is analogous to that of 4.1 and involves the graph of the problem.

Proof. Let R be a rule consisting of more than one action. In the first round, Doer will take one of the actions similar to one of the allowed actions (recall that by definition of S_i, x_i is similar to x_i). Denote that action taken in the first round a.

Note, that $U(a)$ is the highest among all actions Doer could have made in the first round.

By construction of G, there exists a path from v_i to v_j only if $U(x_i) < U(x_j)$. Thus, there do not exist any paths from a to any other actions the doer could have made in round one. This implies, that we can remove them from R with no impact on Doer's decision-making. □

Reminder: S_i is the set of actions similar to x_i, v_i is the vertex corresponding to x_i

Corollary 4.3.1. *For any X, S there exists a rule, which is optimal and consists of exactly one action.*

It is obvious that an optimal rule depends on X and S. However, we can show that it also depends on the number of periods n.

Proposition 4.4. *Optimal rule depends on the number of periods.*

Proof. Consider TF. Changing the number of periods five to four, R_{Naive} (50 minutes of work) becomes better, than $\{f_5\}$ (45 minutes of work). □

5 Imperfect Knowledge of the Set of Similar Actions

Returning to the TF example, it is not clear why S is as presented there. If browsing Twitter for 10 minutes is similar to browsing Twitter for 5 minutes, then wouldn't browsing Twitter for 7 (or 6 or 8, or 4.5) minutes also be similar to browsing Twitter for 5 minutes, i.e. it seems that:

Reminder: S_i is the set of actions similar to x_i, S is the collection of all S_i

$$t_{10} \in S_{t_5} \implies t_7 \in S_{t_5}$$

There are two chief reasons for this presentation. Firstly, mental processes are coarse [25], meaning, we are not very good at distinguishing 0.1 seconds from 0.2 seconds, and in everyday life, neither we are very good at distinguishing 3 minutes from 5 minutes. Once we accept the coarseness of time-perception, the intervals we presented in TF example lead to no loss in generality.

Even so, the TF is example is prohibitively simplistic. In reality, there are dozens if not hundreds of alternative actions available at every turn. Are we always aware of all of them? This does not seem to be true: attention, computation, and memory are all limited resources in the human brain (even more so for our Doer). Thus, we relax the assumption of perfect knowledge of S.

It is plausible that rather than knowing S in advance, Doer discovers its constituents gradually. For example, recall the TF case, but suppose that there are not just Twitter and Facebook, but hundreds of distracting websites. Then, at the beginning of the period, Doer might only be able to think of several of them and choose action from the set of similar actions he was able to think of.

We will augment the model with the mental process of discovery.

Assume that Planner is Sophisticated and has perfect knowledge of S, but Doer doesn't know any part of S, except for allowed actions. Denote the part of S known by Doer as K. At the start of each period, prior to choosing the action to take, Doer randomly discovers 1 action similar to $x \in K$ (if there is such an action).

Proposition 5.1. *In some cases, imperfect knowledge of S will alter Sophisticated Planner's optimal rule.*

Proof. Consider the following case:
$U(w) = 0$
$U(f_5) = 1$
$U(t_{10}) = 2$
$S_w = \{w,\ f_5,\ t_{10}\}$
$S_{f_5} = \{w,\ f_5\, t_{10}$
$S_{t_{10}} = \{w,\ f_5, t_{10}\}$

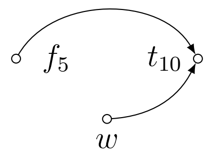

fig. 4

Suppose, there is only one period. Clearly, if Doer knows S perfectly, then all rules are equivalent and Doer will choose action t_{10}, spending 15 minutes working.

However, consider rule w, when Doer has imperfect knowledge of S. Then, with $p = 0.5$ he will discover that f_5 is similar to w and with $p = 0.5$ he will discover that t_{10} is similar to w. In the first case he will choose f_5 and work for 20 minutes, while in the second case he will choose t_{10} and work for 15 minutes. The expected amount of working time is $\frac{20}{2} + \frac{15}{2} = 17.5$. $\square$

This example is instructive in that "semi-sophisticated" Planner who, on the one hand, knows that Doer will not follow the rule precisely, but, on the other hand, believes that Doer knows S perfectly will essentially give up in some cases, believing that the rule she devises will not make any difference.

"Fully sophisticated" Planner knows that sometimes she can set rules that will not necessarily lead Doer to a different equilibrium (like in the original TF case), but will confuse him just enough, so that in this confusion he would get more work done.

In further research we plan to investigate how the number of similar behaviors Doer is able to discover at once impacts the optimal rule.

6 Conclusion

In this paper we examined how people might "slip away" from their commitments to themselves without noticing it. We considered a two selves model, consisting of Planner and Doer, with Planner devising the rule for the long-term benefit of a person and Doer following his perception of these rules. The perception is determined by the similarity among possible actions, with Doer choosing the action he gains most utility from, from the set of actions he perceives to be similar to allowed ones.

Most importantly, we have proved that Sophisticated Planner (who has complete access to Doer's mental processes) will devise the rule strategically, meaning, in some cases, she will deliberately allow some unproductive behaviors, in order to bring under control the amount of time Doer spends unproductively. We have also shown that there always exists a simple optimal rule, consisting of only one action allowed. Further, we have shown that, provided enough periods, regardless of the rule, Doer will settle into a specific action, from which he will not deviate.

We conclude that the key avenue towards increasing are (1) strategic self-restraint: it's preferable to sometimes indulge in unproductive behaviors, rather than try to work constantly; (2) external monitoring: the reason our commitments to ourselves fail is inability step out of the situation and recognize that similarity between behavior does not imply equivalence in whether they are allowed or not. A possible way to deal with this shortcoming is to ask for external monitoring, so that the person observing our actions and judgments could substitute for inattentive internal Planner.

References

[1] R. F. Baumeister and T. F. Heatherton, "Self-Regulation Failure: An Overview," *Psychological Inquiry*, 1996. DOI: 10.1207/s15327965pli0701_1.

[2] K. D. Vohs and R. F. Baumeister, *Handbook of Self-Regulation, Third Edition: Research, Theory, and Applications*. Guilford Publications, 2016, ISBN: 978-1-4625-2045-9.

[3] R. H. Strotz, "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1955. DOI: 10.2307/2295722.

[4] R. H. Thaler and H. M. Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy*, 1981.

[5] G. Loewenstein, "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Processes*, 1996. DOI: 10.1006/obhd.1996.0028.

[6] S. Frederick, G. Loewenstein, and T. O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 2002. DOI: 10.1257/002205102320161311.

[7] R. Bénabou and J. Tirole, "Self-Knowledge and Self-Regulation: An Economic Approach," *The Psychology of Economic Decisions*, 2003.

[8] D. Laibson, "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 1997. DOI: 10.1162/003355397555253.

[9] T. O'Donoghue and M. Rabin, "Doing It Now or Later," *American Economic Review*, 1999. DOI: 10.1257/aer.89.1.103.

[10] ——, "Addiction and self-control," in *Addiction: Entries and exits*, New York, NY, US: Russell Sage Foundation, 1999, ISBN: 978-0-87154-235-9.

[11] F. Gul and W. Pesendorfer, "Temptation and Self-Control," *Econometrica*, 2001.

[12] F. Gul, "The Simple Theory of Temptation and Self-Control†," *Manuscript*, 2005.

[13] A. Etzioni, "The Case for a Multiple-Utility Conception," *Economics and Philosophy*, 1986. DOI: 10.1017/S1478061500002619.

[14] G. F. Loewenstein and T. O'Donoghue, "Animal Spirits: Affective and Deliberative Processes in Economic Behavior," *SSRN Electronic Journal*, 2004. DOI: 10.2139/ssrn.539843.

[15] D. Fudenberg and D. K. Levine, "A Dual-Self Model of Impulse Control," *The American Economic Review*, 2006.

[16] R. Bodner and D. Prelec, "Self-Signaling and Diagnostic Utility In Everyday Decision Making," *The Psychology of Economic Decisions*, 2003.

[17] G. Ainslie, *Breakdown of Will*. Cambridge University Press, 2001.

[18] R. Bénabou and J. Tirole, "Willpower and Personal Rules," *Journal of Political Economy*, 2004. DOI: 10.1086/421167.

[19] J. Benhabib and A. Bisin, "Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption–saving decisions," *Games and Economic Behavior*, 2005. DOI: 10.1016/j.geb.2004.10.004.

[20] N. Mazar, O. Amir, and D. Ariely, "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance," *Journal of Marketing Research*, 2008. DOI: 10.1509/jmkr.45.6.633.

[21] J. L. Freedman and S. C. Fraser, "Compliance without pressure: The foot-in-the-door technique.," *Journal of Personality and Social Psychology*, 1966. DOI: 10.1037/h0023552.

[22] J. M. Burger, "The Foot-in-the-Door Compliance Procedure: A Multiple-Process Analysis and Review," *Personality and Social Psychology Review*, 1999. DOI: 10.1207/s15327957pspr0304_2.

[23] I. Brocas and J. D. Carrillo, "The Brain as a Hierarchical Organization," *American Economic Review*, 2008. DOI: 10.1257/aer.98.4.1312.

[24] A. Tversky, "Features of similarity," *Psychological Review*, 1977. DOI: 10.1037/0033-295X.84.4.327.

[25] S. Mullainathan, J. Schwartzstein, and A. Shleifer, "Coarse Thinking and Persuasion," *The Quarterly Journal of Economics*, 2008.

[26] I. Gilboa and D. Schmeidler, "Case-Based Decision Theory," *The Quarterly Journal of Economics*, 1995. DOI: 10.2307/2946694.