THE USE OF ROBO-ADVISORS AS COMMITMENT DEVICES

Alexia Gaudeul
Caterina Giannetti

THINK FORWARD INITIATIVE

THE USE OF ROBO-ADVISORS AS COMMITMENT DEVICES [*]

TECHNICAL REPORT

Alexia GAUDEUL[1] and Caterina GIANNETTI[2]

December 2020

Abstract

We report results from a stock trading experiment in which we tested the use of commitment devices to help individuals remedy with the disposition effect – the tendency to sell winning stocks and keep those that lose value. A commitment device is a self-imposed arrangement that helps people to stick to a decision they have made. Our experiment lasted three weeks. In the first week, participants could trade without any type of restrictions. In the second week, participants had to rely on algorithms (robo-advisors) every two periods. One type of adviser simply blocked trading every two rounds. Another type traded for them every two rounds according to Bayesian rules. We also varied the rigidity of both types of algorithm, whereby some participants could override algorithmic trading, and some could not. Finally, in the last week, participants decided whether to use the robo-adviser they had had to use in the second week. The majority of participants chose not to rely on an algorithm, but still, adoption was higher for robo-advisers that were doing active trading and could be overridden. Participants who performed better on their own – those with lower levels of the disposition effect – were more likely to opt for a robo-adviser in the third week. Allowing overrides of advice did not significantly reduce trading performance. This shows that those who are more subject to the disposition effect are less likely to seek advice, this issue can be alleviated by offering active robo-advisers that can be overridden.

Keywords: disposition effect, commitment devices, robo-advisors, sophisticated investors, trading

---

[1] Alexia Gaudeul, University of Göttingen, Email address: alexia.gaudeul@wiwi.uni-goettingen.de
[2] Caterina Giannetti, University of Pisa, Email address: caterina.giannetti@unipi.it

Introduction

The use of artificial intelligence in guiding the activity of humans is an exciting, promising but also slightly frightening new area of research. Already, increasingly-clever applications help us manage our time, make decisions and keep our commitments, to ourselves and to others. In particular, new technologies are revolutionizing the way we seek and receive advice. For example, applications make use of artificial intelligence to help customers correct their eating habits (e.g. Weight Watchers App), choose a suitable dating mate (e.g. through partnership websites) and adjust their sports training strategies (e.g. Nike+ Running App).

This is also the case for **financial advice**, where **robo-advisors** that provide automated and personalized portfolio management advice have rapidly entered the field. Robo-advisors can help people correct for the impact of irrational factors in their decisions. They are particularly attractive because of their low cost, permanent availability and easy access via user-friendly interfaces. Robo-advisors also hold the potential to reduce moral hazard problems in the relationship between advisor and the investor, as they can be verifiably designed to unambiguously serve the interests of the investor rather than those of the advisor.

In this project, we set up a series of experiments to study how robo-advisors can be better designed to help investors commit to rational investment strategies. In particular we aim to study how different types of algorithms (i.e. robo-advisors) can be used to reduce the **disposition effect,** i.e. the individual tendency to sell winning stocks and keep losing stocks.

The **disposition effect** is one of the most studied, and well documented of the various market trading anomalies. In particular, this type of behavior leads to portfolios that are over-weighted in loss-positions, thus reducing investor performance. Furthermore,

neural tests show that individuals experience regret due to this behaviour, thus indicating that their decisions are ex-post sub-optimal even from their own point of view (Frydman and Camerer 2016). While the disposition effect is a robust and well documented empirical phenomenon, less attention has been devoted to helping individuals cope with it.

We therefore investigate how individuals who are subject to the disposition effect may paradoxically attain better outcomes by reducing their own freedom of action, that is, by committing either to trade less often or by letting an algorithm decide how to trade.

Our objective is not only to understand how robo-advisors can be designed to improve people's decisions – that is indeed a low bar to pass! – but also which features make them more attractive for people to adopt. The issue indeed is that investors may not be convinced at first that such algorithms can be useful.

In order to improve acceptance, we therefore vary across treatments how strongly people may commit to the use of the algorithm. Our proposed algorithm was thus combined with a **commitment device** in some treatments, whereby traders had to commit to using that algorithm. We thus consider if being able to override the algorithm makes people more likely to accept it. By varying the degree of commitment imposed by the algorithm (soft, with override, vs. hard, with no override), we aim to identify the optimal design of robo-advisors so as to increases the use of robo-advisors (their *take-up rate*), and thus reduce the disposition effect.

We consider if commitment improves their financial decisions, that is, whether traders performed better with hard commitment. In cases where the adviser could be overridden, we also investigate whether individuals do override. Finally, we consider whether bad traders have

sufficient insight into their issue to accept such restrictions on their ability to trade.

Our results suggest that participants can achieve better performance by restricting their ability to trade. However, only a small minority of participants is apparently aware of this benefit (i.e. are sophisticated) and decide to rely on algorithms after having tried them. The large majority of our participants prefer to avoid any type of constraints on their own behaviour, especially those who would benefit the most from such constraints. Participants that relied on algorithms performed significantly better than participants that decided to play freely, both in terms of disposition effect and earnings. Encouragingly however, we observe larger take-up rates for soft algorithms, i.e. the ones that participants are able to override, and low rates of overriding in this case.

Background

We build our research on previous experimental settings for trading in artificial stock markets (e.g. Frydman et al 2014, Frydman and Rangel 2014). Participants in our experiment were invited to trade in such an artificial stock market, via a link that was accessible on their device (either computer or mobile phones), over a period of three weeks. Participants were paid based on the value of their portfolio at the end of the experiment: the higher its value, the more they earned.

1.   The disposition effect

Several researchers have already studied the disposition effect, both empirically and experimentally. A variety of theories have also been proposed to explain it: prospect theory (e.g. Ling and Yang 2013); regret minimisation (e.g. Bleichrodt et al 2010); realisation utility (e.g Frydman et al 2014, Barberis and Xiong 2009). While the underlying causes of the disposition effect is still debated, the evidence on this phenomenon is extremely robust.

In particular, household investors are more affected by the disposition effect than professional investors, and the disposition effect is greater for females and older people (Dhar and Zhu, 2006; Rau, 2014). National culture also seems to play a role. Indeed, populations that are more focused on the long-term and more willing to reject strict social norms have lower average levels of the disposition effect (Breitmayer 2019).

To measure the disposition effect, researchers usually refer to the original work of Odean (1998) and assign each stock to four categories:
- "***realised gain***": a stock that is sold at a price that is higher than the purchasing price;
- "***paper gain***": a stock that is *not* sold but whose price is higher than the purchasing price;
- "***realised loss***": a stock that is sold at a price that is lower than the purchasing price;

- "***paper loss***": a stock that is *not* sold but whose price is lower than the purchasing price;

The disposition effect is computed as DIFF, the difference between the proportion of realised gains (PGR) and losses (PLR), that is:

$$Diff = \frac{\#RealizedGains}{\#(RealizedGains + PaperGains)} - \frac{\#RealizedLoss}{\#(RealizedLoss + PaperLoss)} \quad (1)$$

In this research we additionally compute DIFF_AMOUNT, which takes account of the magnitude of the gains and losses:

$$Diff\_Amount = \frac{ECURealizedGains}{ECURealizedGains + ECUPaperGains} -$$

$$- \frac{ECURealizedLoss}{ECURealizedLoss + ECUPaperLoss} \quad (2)$$

Both indicators have a theoretical range going from -1 to +1, where +1 is the value for an investor that sells all his winning positions and holds all losing ones, -1 is the value for an investor that sells all losing positions and holds all winning ones, and 0 is the value for an investor who behaves the same in both cases. The higher the values of these indicators, the more an individual is subject to the disposition effect. In our experiment, the optimal strategy results in both indicators being negative. As a result, a positive value for these indicators unequivocally identifies a distorted – i.e. non optimal - behaviour.

While the disposition effect is well documented, very little attention has been devoted to mechanisms that would allow investors to cope with it. Frydman and Rangel (2014) show that it is possible to reduce the disposition effect by decreasing the saliency with which the purchasing price is disclosed. Fischbacher et al (2017) show that giving subjects access to automatic

selling devices (stop-loss and take-gain orders) increases the proportion of realized losses.

2.   Soft vs. hard commitment device

As stated above, one of the main ideas behind this research is to examine whether traders are ready to use robo-advisors that act as commitment devices, i.e. force people to stick to decisions they have made. We also consider if such strong commitment to the advisor is necessary to derive benefits from advice.

Individuals may freely choose to limit their freedom of choice by committing to follow the advisor. This is a rational decision if individuals, aware that the advisor is on average better, but also aware that they might be tempted to override the advisor, decide to limit their own freedom to do so (Bryan, Karlan, and Nelson 2010). This implies however a high level of rationality.

One way robo-advisors may commit people is by simply restricting future choices, another is softer and involves financial or psychological incentives. For example, online platforms like StickK allow people to bet financial and reputational stakes on their commitments. In that case, people can still override their commitment, but at a cost.

This is the distinction that Bryan et al. (2010) make between hard and soft commitment devices: if the incentives (rewards or avoidance of costs) are primarily financial, it is called a hard device; if they are mostly psychological, it is a soft device.

In general, soft commitments are more likely to be adopted, but do not work as well as hard commitments. For example, Fischbacher et al, 2017, showed that simple reminders (i.e. reminders of price limits at which participants wanted to sell the assets) did not reduce the disposition effects.

Sophisticated agents would opt for hard commitments, but less sophisticated ones, who would reject hard commitments, may still accept softer commitments (Beshears et al., 2018; Dupas and Robinson, 2013; Royer et al., 2015; Beshears et al., 2015; Burke et al., 2018; Duckworth et al., 2016; Bryan et al., 2010).

In our experiment, we therefore vary whether traders commit to following the robo-advisor, or can override its decision. A robo-advisor that can be overridden can be likened to a nudge (Thaler & Sunstein, 2008). This implies a potential loss in efficiency if the advisor is more likely to be right than the trader but is less threatening to individual freedom of choice.

Experimental Design

The set-up of our experiment closely resembles Frydman et al (2014). However, we introduce an important change by moving the experiment from the physical lab, where participants stay for only a limited time in a specific room, to a different setup, where participants take part in the experiment while going on with their usual activity. We programmed our experiment with *oTree* (Chen et al 2016), which allowed us to put the experiment online so that each one of our participants had a personalised weblink to play. We relied on *GMass* to send a personalised email to participants every 8 hours to remind them that a new market session had begun.

This move outside the physical laboratory allows us to increase the interval between trading periods from seconds to hours (and days). This is important as financial choices are the outcome of the interaction between an instinctive-affective mechanism (System 1) and a deliberative-cognitive mechanism (System 2) (Kahneman 2002, Hirshleifer 2015). By extending the time intervals between consecutive rounds, we allow participants to develop their understanding of the market and devise their own strategies in a more reflective fashion than they have the time for in the usual experiment. Our design thus allows us to further test the robustness of the disposition effect to experience.

Another difference from Frydman and Rangel (2014) is that our participants could buy and sell simultaneously up to 3 stocks per round, while in their setting participants could buy and sell only one stock at time. Specifically, in each round each one of the three stocks (A, B, C) had its price randomly updated. The price path of each stock was independently governed by a two-state Markov chain, with a good state and a bad state. If the stock *i* is in the good state, its price increases with probability 0.70 while it decreases with probability 0.30. If the stock is instead in a bad state, its price increases

with probability 0.30 and it decreases with probability 0.70. Independently of the direction of the price change, the magnitude of the price variation is uniformly drawn from {5, 10, 15}. In subsequent rounds, the (good or bad) state of each stock remains the same with probabilities 0.80, while it switches state with probability 0.20. To make comparisons easier across participants and treatments, we predetermined 6 series of price realisations, the same across treatments.

Each subject could hold a maximum of one share of each stock and a minimum of zero (i.e. short-selling was not allowed). The trading decision was therefore reduced to deciding whether to sell a stock (conditional on holding it) or buying a stock (conditional on not holding it). As in Frydman and Rangel (2014), and Frydman et al (2014), each stock exhibits positive autocorrelation. In other words, a stock that performed well in the last round is likely to be in a good state in the subsequent round.

The experiment lasted 21 days, with each day having 3 trading rounds, each lasting 8 hours. In the first 7 days (i.e. the first week), participants played the base-game without an algorithm to support their choice (i.e. without any type of robo-advisors). In the following 7 days (i.e. the second week), a new start was made and participants had access to an advisor.

Algorithms differed in type (i.e. either Blocked trading, or Bayesian trading) and in the flexibility of the commitment to use them (i.e. soft vs hard commitment). Specifically, the blocked trading algorithm committed participants to trade only every two rounds, while the Bayesian algorithm traded according to the Bayesian updating of the probability of the stock being in a good state. If the algorithm was of a soft-type, participants were always free to override the algorithm trading choice, while they could not do so if the algorithm was of the hard-type.
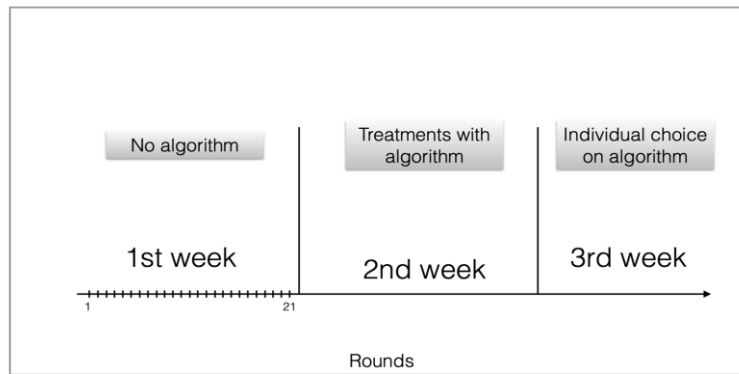
Figure 1: Timeline, algorithm (i.e robo-advisor) imposed in the second week, optional in the third week

The algorithm that prevented trading was meant to help people make more considered decisions, while the one that made trading decisions was meant to help them learn how to trade optimally. Making the advice optional, that is, not forcing individuals to follow it, was meant to overcome algorithm aversion, that is, let people change decisions made by the algorithm so they feel more in control.

In the third week, a new start was again made and participants selected their favourite way to play, that is, either as in the first or as in the second week - i.e. with or without an algorithm (see experiment structure in Fig. A and Instructions at the end of the report). To summarise, we implemented the following treatments:

1. **Hard Blocked**: participants have to let an algorithm (advisor) make decisions every two rounds, and this decision is not to trade;

2. **Soft Blocked**: participants rely on the same algorithm as above, but they can override that decision on a period by period basis;

3. **Hard Bayes:** participants have to let an algorithm trade every two rounds according to a Bayesian updating of probability, i.e. sell/not buy a stock whenever the probability of that stock being in a bad state is above 50% (and vice-versa);

4. **Soft Bayes:** participants rely on the same algorithm as above, but they can override that decision on a period by period basis;

To control for possible order effect, we run four additional **reverted** treatments, in which in the first week participants play with one of the four advisors while in the second week they play freely. As in the standard treatments, in the last 7 days (i.e. third week) participants needed to choose whether they preferred to play as in the second week or as in the first week. We denote those treatments as **Hard Blocked reverted, Soft Blocked reverted, Hard Bayes reverted** and **Soft Bayes reverted.**

Participants were paid the value of their portfolio at the end of one of the 3 week selected at random.

See Figure 1 for a timeline in the case where the use of the algorithm (i.e. robo-advisor) is imposed in the second week.

1.   Theoretical predictions

As in the paper of Frydman et al (2014), and Frydman and Rangel (2016), our set-up induces positive autocorrelation in stock price changes, which implies that a risk-neutral rational trader ought to sell losing stocks more often than winning stocks, thereby exhibiting the *opposite* of the disposition effect. In particular, the optimal trading strategy for a subject is to sell (or not to buy) a stock when he believes that it is more likely to be in a bad state than in a good state, and to buy (or hold) a stock when he believes that it is more

8

likely to be in good state. Since the three stocks are uncorrelated in our experiment, it is rational for the participants to consider each stock individually.

We can define optimal trading more precisely. Let p_{it} be the price of stock i in round t and let $q_{it} = Pr(p_{it}, p_{i,t-1}, \ldots . p_{i1})$ be the probability, from the point of view of a rational (Bayesian) investor, that stock i is in the good state. Let $z_{it} = 1$ indicates a price increase for the stock i, and $z_{it} = -1$ indicates a price decrease. Then, we have

$$q_{i,t}(q_{i,t-1}, z_{it}) = \frac{Pr(s_{it} = good)Pr(q_{it})}{Pr(s_{it} = good)Pr(q_{it}) + Pr(s_{it} = bad)Pr(q_{it})}$$

$$= \frac{(0.5 + 0.2z_{it})(0.8q_{i,t-1} + 0.2(1 - q_{i,t-1}))}{(0.5 + 0.2z_{it})(0.8q_{i,t-1} + 0.2(1 - q_{i,t-1})) + (0.5 - 0.2z_{it})(0.8(1 - q_{i,t-1}) + 0.2q_{i,t-1})}$$

The optimal strategy is to sell (if holding) or not to buy (if not holding) a stock i when $q_{i,t} < 0.5$, and to keep (if holding) or buy (if not holding) otherwise. The strategy of the Bayes adviser is based on this probability.

Similar to previous experiments (Frydman et al 2014, Frydman and Ragel 2016), it is difficult for participants to exactly compute this probability. However, it is possible to approximate this optimal strategy with a simple rule of thumb: *i.e. "hold on stocks that have recently performed well, sell stocks that have recently performed poorly".*

2.   Participants and experimental protocol

Our first experimental sessions started in October-November 2019. We conducted several pilot sessions until May 2020 (about 120 participants) to refine the setting and design of our experiment. During this period, we also experienced the Covid-19 lockdown restrictions. This did not affect our data collection but forced us to move all the phases of our experiment entirely online. The final data collection started in June 2020 (right after the end of lockdown restrictions in Italy) and ended in Mid-August. Therefore, all our data was collected while the situation in Italy was quite stable.

Participants were randomly selected out of a pool of about 3000 students from 20 departments of the University of Pisa. Participants were invited to the LES laboratory online with *Microsoft Teams* where they received instructions and a **personal weblink** they could use to play (either on their computer or on their mobile phone). There was a trial session with two periods of trading and participants could ask for clarifications. Although always available online, instructions were also read aloud on *Teams* during the explanatory session. Participants also received instructions in a PDF version (see the English Translation at the end of the report). In addition to having a general description of the experimental market, participants were told that at the end of the second and third week they would receive additional information about variations that would be introduced in the game. In addition, every 8 hours participants received an email reminding them the beginning of a new trading session, as well as their personal link to play.

They were finally told that, at the end of the third week, one randomly selected week would be selected for payment. At this time, they needed to write an email to the experimenter in order to receive their payment by bank transfer.

Results

As stated above, data collection started in June 2020 (right after the end of lockdown restrictions in Italy) and ended in Mid-August. Slightly more than 450 participants, mainly students from the University of Pisa, took part in the experiments. Fig. 2 shows the fields of study of our participants, with a large majority of engineers and economists.
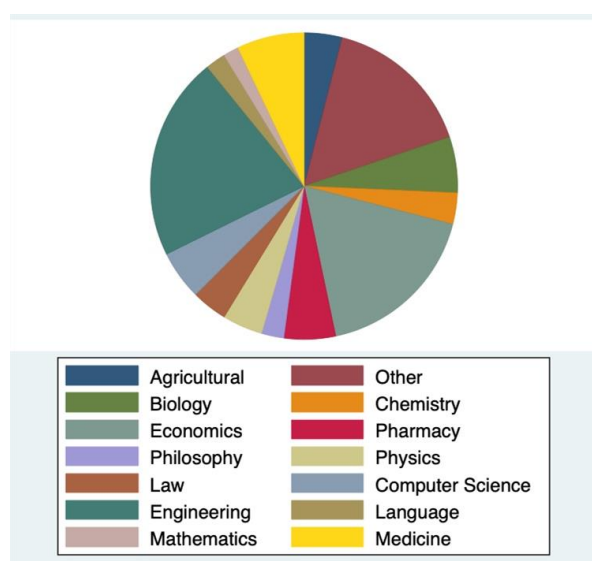

Fig 2: Participants' field of studies

The average payment for participation was about 17.80 Euro, including a show-up fee of 5 Euro.

The participation rate was quite high in the experiment, with a low dropout rate (about 6%) resulting in a sample of 422 participants who went through all phases and claimed payment at the end of experiment. We deleted from our sample 4 participants who did not actively play at any time during the experiment but claimed payment.

Conditional on being in the sample, participants' activity rate was quite high and stable during all three weeks (see Table 1). Participants trade 56% to 73% of periods,

meaning about twice per day out of three possible periods. This share is lower, around 40%, in some Blocked treatments. However, there are no significant differences across treatments in terms of participation rates on average.

Share of weekly activity per treatments				
	1st week	2nd week	3rd week	# obs (participants)
Hard Blocked	0.78	0.69	0.65	58
Hard Blocked reverted	0.76	0.72	0.69	49
Hard Bayes	0.73	0.50	0.51	53
Hard Bayes reverted	0.67	0.58	0.57	47
Soft Blocked	0.75	0.59	0.55	57
Soft Blocked reverted	0.73	0.64	0.55	51
Soft Bayes	0.65	0.51	0.37	49
Soft Bayest reverted	0.76	0.50	0.59	54
Tot Soft	0.72	0.56	0.52	211
Tot Hard	0.74	0.62	0.60	207
Total	0.73	0.59	0.56	418

Table 1: Weekly rate of activity and number of participants across treatments

At the beginning of the experiment, we also collected information about participants' cognitive ability and level of concern for the future, as well as their financial literacy, locus of control and risk-aversion. Results are shown in Table 2. There are no significant differences across treatments. On average, participants were able to answer correctly two out three logical questions (CRT,

Frederick, 2005), slightly more than two out of three basic financial questions, and scored about 37 (min 0, max 94) in the consideration of future of consequence scale (CFC, 12-item scale developed by Strathman et al. 1994). Participants rated on average on a 4-likert scale of risk-aversion. The average participants' age is about 25 years old, and 46% of participants are male.

Participants' characteristics		
	Mean	Sd
CRT	2.02	(1.08)
Future Attitude (CFC)	36.74	(4.41)
Financial Literacy	2.39	(0.72)
Locus of Control	8.22	(1.74)
Risk Aversion	2.26	(0.82)
Age	24.97	(3.99)
Male	46%	

Table 2: Participants' characteristics

1. Do people suffer from the disposition effect?
As stated above, our setting implies that an expected value maximiser would exhibit the opposite of the disposition effect -- that is, negative DIFF values. In Table (3) we report the average disposition effect for our participants. For comparability across treatments, we rely only on choices made *without* the help of a robo-advisor (i.e. excluding choice made by the algorithm). We compute the average disposition effect by looking at both the number and value of stocks sold and bought (DIFF and DIFF_AMOUNT, cf. equations 1 and 2).

As Table (3) highlights, the value of DIFF and DIFF_AMOUNT was relatively high, positive and statistically significantly more than zero in the first week in almost all treatments, though less so in reverted treatments, where participants were helped every two periods by an advisor. In the second week, the disposition effect was lower in almost all treatments. Indeed, the difference between the first and second week (i.e. the Δ column in Table 3) is always positive and significant in some cases. One reason for this

reduction may be that participants learned over time to trade better.

Treatment	Week 1	Week 2	Δ
Hard Blocked	0.146***	0.145***	0.001
Hard Blocked reverted	-0.037	-0.142**	0.104*
Hard Bayes	0.091*	0.028	0.063
Hard Bayes reverted	0.144**	-0.056	0.200**
Soft Blocked	0.159**	0.111	0.048
Soft Blocked reverted	0.014	-0.079	0.065
Soft Bayes	0.131***	-0.011	0.142**
Soft Bayes reverted	0.076	0.021	0.055

Table 3: Average disposition effect: each cell includes the indicator DIFF (eq 1) computed relying only on individual choices, i.e. without considering the choice of the algorithm.

Treatment	Week 1	Week 2	Δ
Hard Blocked	0.175***	0.156***	0.019
Hard Blocked reverted	-0.031	-0.124**	0.093
Hard Bayes	0.130**	0.029	0.101
Hard Bayes reverted	0.157**	-0.050	0.207**
Soft Blocked	0.197***	0.116*	0.081
Soft Blocked reverted	-0.013	-0.104	0.091
Soft Bayes	0.154***	0.014	0.140*
Soft Bayes reverted	0.096*	0.027	0.069

Table 4: Average disposition effect: each cell includes the indicator *Diff Amount* (e. g. 2) computed relying only on individual choices, i.e. without considering the choice of the algorithm.

2. Do people adopt algorithms?

In the last week of our experiment, participants could decide in which way to play the remaining rounds of the game, i.e. whether to play with the assistance of an algorithm or not.

In that situation, we know from the experimental literature that at least a fraction of individuals will behave in a sophisticated way. These individuals are not free from present bias (i.e. tendency to want to experience benefits right away and to postpone the realisation of losses as much as possible) but they are aware this bias hurts them. Therefore, they will be willing to use commitment devices so as to impose the behaviour they planned now for the future on their "future self".

Not all subjects are that sophisticated though. Most subjects display *naïve* behaviour – they are subject to the disposition effect and unaware of being so. As a result, they will not want to make use of any commitment device.

Indeed, on average, the take-up rate is quite low, as only about 36% of our participants decided to rely on a trading algorithm (see Table 5). Importantly, we observe that in the majority of treatments participants were less likely to opt for an algorithm in the third week if they suffered from the disposition effect in the first two weeks (see Table 5). On average, among those who opted to rely on an algorithm in the third week, only 45% suffered from the disposition effect, while among those who decided *not* to opt for an algorithm this share is about 54%.

Treatments	Choice	Share of people with DE	Take-up rate
Hard Blocked	No Algorithm	62%	19%
	Algorithm	63%	
Hard Blocked reverted	No Algorithm	45%	37%
	Algorithm	33%	
Hard Bayes	No Algorithm	41%	40%
	Algorithm	43%	
Hard Bayes reverted	No Algorithm	40%	32%
	Algorithm	35%	
Soft Blocked	No Algorithm	72%	43%
	Algorithm	50%	
Soft Blocked reverted	No Algorithm	59%	49%
	Algorithm	45%	
Soft Bayes	No Algorithm	60%	37%
	Algorithm	65%	
Soft Bayes reverted	No Algorithm	44%	36%
	Algorithm	25%	
Total	No Algorithm	54%	36%
	Algorithm	45%	

Table 5: Share of people suffering from the disposition effect (DE) in the first/second week and algorithm choice in the third week choice.

	algorithm (1)	algorithm (2)	algorithm (3)	algorithm (4)
DE	-0.076**	-0.076**	-0.074**	-0.070**
Reverted	0.028	0.028	0.029	0.016
Soft	0.076*	0.076*	0.070	0.070
Bayes	0.070	0.070	0.063	0.063
mean activity		0.001	0.014	-0.011
CRT			0.012	0.012
CFC			-0.010*	-0.009
Financial Literacy			-0.057	-0.053
Control				-0.006
Risk aversion				0.077***
Obs	417	417	417	417

Table 6: Determinants of the algorithm take-up. The dependent variable is a dummy equal 1 if the individual opted for an algorithm in week 3 (and zero otherwise)

We get a synthetic view of the effect of each variable on the take-up rate by running a logistic regression of the type

$$Logit(Algorithm) = \beta_1 DE + \beta_2 Reverted + \beta_3 Bayes + \beta_4 Soft + \\ + \beta_5 Mean\ Activity + \beta_6 CRT + beta_7 CFC \\ + \beta_8 Financial + \beta_9 Control + \beta_{10} Risk\ aversion + u$$

where the dependent variable is the dummy *Algorithm* equal to 1 if an individual opted for the algorithm in the third week (and zero otherwise). *Soft* is a dummy variable equal to 1 if an individual could opt for a soft algorithm (and zero otherwise). *Reverted* is a dummy variable equal to 1 if the individual played with an algorithm in the first week (and zero otherwise). *Bayes* is a dummy equal to one if the algorithm was Bayesian (and zero otherwise). *DE* is a dummy variable equal to 1 if the individual suffered from the disposition effect in the week he could freely play (i.e. Diff>0). We also control for the level of individual average level of activity during the first two weeks (*mean activity)* and other individual characteristics (such as CRT score, financial literacy, CFC score, locus of control and risk aversion). Results are reported in Table 6. In line with what was observed in Table 5, in column (a) we observe that being subject to the disposition effect decreases the probability of taking up the algorithm by 7.6%, and this is economically and statistically significant.

Although soft algorithms seem to be preferred (+7.6%) the effect is only marginally statistically significant. The more sophisticated type of algorithm (i.e. *Bayes)* is also preferred (+7.0%) although the effect is again not statistically significant.

The level of weekly activity (column 2) does not affect the likelihood that a participant will take-up an algorithm in the third week. Among individual characteristics (column 3), the attitude towards future consequences (CFC) plays a marginal role, decreasing the probability to take-up an algorithm by 1% for an increase by 1 in the scale, along with risk aversion, which instead increases significantly the probability to take-up the algorithm by 7.7% for an increase of 1 point in the scale of risk-aversion.

3. Do people benefit from the advisor?
Finally, we are interested in whether individuals who decided to rely on an algorithm in the end performed better than individuals who did not.

Thus, in a last regression, we regress various indicators of the strength of the disposition effects in each week on a dummy *Algorithm* equal to one if the individual decided or had to rely on an algorithm (and 0 otherwise). The dummy *Soft* is equal to one if the individual could override the algorithm (and 0 otherwise). Dummy *Bayes* denotes the type of advisor, equal to 1 if the algorithm was Bayesian. Dummy *Reverted* equals one if the individual was in reverted treatments (and 0 otherwise). We also include the same set of individual characteristics as before, i.e. *financial literacy*, future attitude (*CFC*) and cognitive Ability (*CRT*), locus of control (*Control)* and *risk aversion*, as well as individual engagement in the experiment during the week (*Weekly share activity*).

We thus run the following type of regressions:

$$Diff = \beta_1 Algorithm + \beta_2 Reverted + \beta_3 Bayes + \\ + \beta_4 Soft + \beta_5 WeeklyActivity + \beta_6 CRT + beta_7 CFC \\ + \beta_8 Financial + \beta_9 Control + \beta_{10} Risk\,Aversion + u$$

Table (7) - column (1) - highlights that individuals relying on an algorithm had lower levels in their disposition effect than individuals not relying on an algorithm (-0.179, p-value=0.000). This difference is not only statistically but also economically significant, considering the average level of the disposition effect in the sample (about -0.057). Being in the reverted treatments also decreases the level of the disposition effect by -0.093.

Although participants seem to prefer soft algorithms, we do not find any benefit of using them compared to hard algorithms. On the contrary, we see that having the possibility to override an algorithm choice, lead to an increase of 0.038 in the level of the disposition effect. However, relying on a Bayesian algorithm, which is more sophisticated than the simpler restricted trading, was associated with a lower disposition effect (about -0.123, p-value=0.000).

Finally, if we look at participants' characteristics, we observe positive differences (in terms of reduction of the disposition effect) due to weekly rate of activity and other individual traits (such as cognitive ability). In particular, participants who were more engaged in the experiment had a lower disposition effect. A 10% higher participation rate was associated with a -0.0204 lower disposition effect. Similarly, 1 more correct answer in the cognitive reflection test (CRT) is associated with 0.029 lower disposition effect. Financial literacy does not have an impact, neither does having a greater concern for future consequences (CFC scale). Results are analogous if we look at our second indicators Diff Amount, Table (8) - column (2).

A last point to examine is whether participants who relied on an algorithm performed better also in terms of earnings. Thus, in Table (8) - column (3) - we repeat the previous regressions but with earnings (i.e. ECU) as the dependent variable and on the same sets of regressors.

Indeed, we observe that participants relying on a Bayes algorithm earned significantly more (about 12 ECU more) than participants who opted not to rely on an algorithm. Being risk-averse also reduces significantly the level of earnings by about 3 ECU for one-point increase in risk aversion. Other individual characteristics however appear not to have a significant impact.

	Diff	Diff Amount	Earning
ALGORITHM	-0.179***	-0.210***	-0.800
Reverted	-0.093***	-0.106***	-0.446
Bayes	-0.123***	-0.146***	12.002**
Soft	0.024	0.036	-1.189
Weekly activity	-0.204***	-0.224***	-3.493
CRT	-0.029*	-0.030*	2.212
CFC	0.005	0.004	-0.082
Financial Literacy	-0.015	-0.022	-4.194*
Control	0.002	0.001	-0.706
Risk aversion	-0.009	-0.009	-3.425**
2nd week	-0.096***	-0.111***	-45.998***
3rd week	-0.108***	-0.116***	-11.872**
Constant	0.207	0.287	407.274***
Observations	1254	1254	1226

Tab 8. Individual performance over weeks and earnings

Discussion

The results from our research shed light on how individuals may be helped in coping with the disposition effect. In particular, our experimental analysis clearly highlights that there are – as expected – two types of investors: sophisticated and naïve.

The first category of investors is smaller in number and comprises those individuals who are subject to the disposition effect but appear to be aware of it, whereby they adopt measures to combat it. They are willing to restrict their freedom to trade in order to achieve better outcomes. The second category is more numerous and comprises those investors who do not realise that relying on a simple commitment device that let them make decisions only in a restricted set of periods (as implemented in the hard-commitment treatment) would allow them to improve their performance. We also find that individuals prefer more active and complex robo-advisors, that trade for them rather than simply not doing anything. They also prefer soft commitment devices, that is, those that can be overridden. In this sense, leaving the possibility to the individual to override the algorithm encourages take-up by those who are not fully convinced of the benefit of restrictions. Nevertheless, we need to acknowledge that we need to test whether the net effect of increased take-up of soft commitment combined with the lower effectiveness of such commitments is positive, especially for those who need self-commitment the most -- the worst traders.

However, the results from this research, together with evidence emerging from previous related studies, already suggest directions for the design of commitment devices in order to curb the negative consequences of the disposition effect for portfolio management. In particular, two important implications follow to increase take-up rate and overcome algorithm aversion. First, it is crucial to let people experience markets on their own

before offering robo-advising. We found that even bad investors who had the opportunity to use an advisor for a while did not seem to realize how much better they performed with it. Robo-advising should therefore focus on opening the eyes of those individual investors who perform worst on the stock market, so that they realize how much they could benefit from the use of an advisor. Second, it is also important to give people the ability to override the advisor. Simply offering the option to not approve trades by the robo-advisor would be a simple way to enhance the feeling of being in control of decisions, and thus overcome algorithm aversion.

On top of that, we also see interesting avenues for future research. Specifically, it could be of interest to let people have more leeway in the design of their advisors (e.g. by determining by themselves the set-up of the algorithm). For example, it would be interesting to let them vary the strength of their commitment to follow the advisor, e.g. by putting a price on overrides.
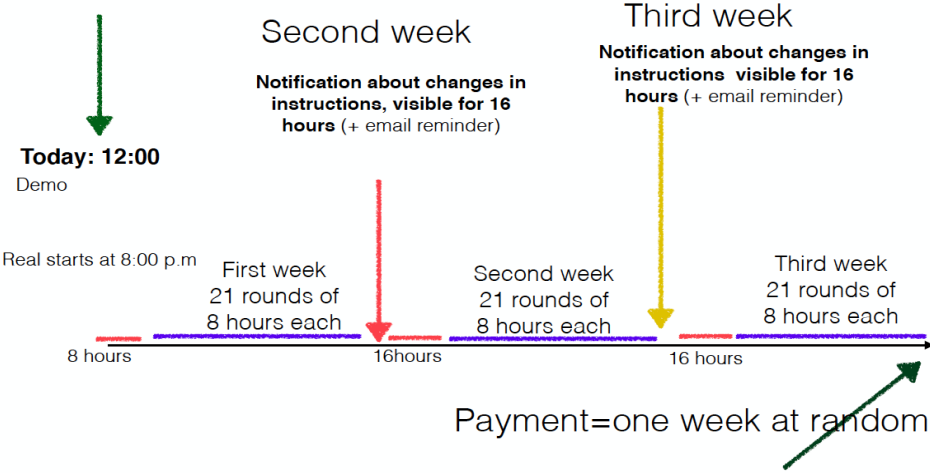
References

Barberis, Nicholas, and Wei Xiong. "What drives the disposition effect? An analysis of a long-standing preference-based explanation." *the Journal of Finance* 64.2 (2009): 751-784.

Beshears, John, et al. "Potential vs. realized savings under automatic enrollment." *Retrieved from the TIAA Institute web site: https://www. tiaainstitute.org/sites/default/files/presentations/2018-07/Potential%20vs%20Realized%20Savings_Beshears_July* 202018 (2018).

Beshears, John, et al. *Self control and commitment: can decreasing the liquidity of a savings account increase deposits?.* No. w21474. National Bureau of Economic Research, 2015.

Braga, Robson, and Luiz Paulo Lopes Fávero. "Disposition effect and tolerance to losses in stock investment decisions: An Experimental Study." *Journal of Behavioral Finance* 18.3 (2017): 271-280.

Breitmayer, Bastian, Tim Hasso, and Matthias Pelster. "Culture and the disposition effect." *Economics Letters* 184 (2019): 108653.

Bryan, Gharad, Dean Karlan, and Scott Nelson. "Commitment devices." *Annu. Rev. Econ.* 2.1 (2010): 671-698.

Burke, Jeremy, Jill Luoto, and Francisco Perez-Arce. "Soft versus Hard Commitments: A Test on Savings Behaviors." *Journal of Consumer Affairs* 52.3 (2018): 733-745.

Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. "Fight or flight? Portfolio rebalancing by individual investors." *The Quarterly journal of economics* 124.1 (2009): 301-348.

Chen, Daniel L., Martin Schonger, and Chris Wickens. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance* 9 (2016): 88-97.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi. "The promises and pitfalls of robo-advising." *The Review of Financial Studies* 32.5 (2019): 1983-2020.

Dhar, Ravi, and Ning Zhu. "Up close and personal: Investor sophistication and the disposition effect." *Management Science* 52.5 (2006): 726-740.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. "Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them." *Management Science* 64.3 (2016): 1155-1170.

Dupas, Pascaline, and Jonathan Robinson. "Why don't the poor save more? Evidence from health savings experiments." *American Economic Review* 103.4 (2013): 1138-71.

Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy. "The causal effect of stop-loss and take-gain orders on the disposition effect." *The Review of Financial Studies* 30.6 (2017): 2110-2129.

Frydman, Cary, and Colin Camerer. "Neural evidence of regret and its implications for investor behavior." *The Review of Financial Studies* 29.11 (2016): 3108-3139.

Frydman, Cary, and Antonio Rangel. "Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price." *Journal of economic behavior & organization* 107 (2014): 541-552.

Frydman, C., Barberis, N., Camerer, C., Bossaerts, P. and Rangel, A., 2014. Using neural data to test a theory of investor behavior: An application to realization utility. *The Journal of finance*, *69*(2), pp.907-946.

Hirshleifer, David. "Behavioral finance." *Annual Review of Financial Economics* 7 (2015): 133-159.

Imas, Alex. "The realization effect: Risk-taking after realized versus paper losses." *American Economic Review* 106.8 (2016): 2086-2109.

Li, Yan, and Liyan Yang. "Prospect theory, the disposition effect, and asset prices." *Journal of Financial Economics* 107.3 (2013): 715-739.

Ploner, Matteo. "Hold on to it? An experimental analysis of the disposition effect." *Judgment & Decision Making* 12.2 (2017).

Royer, Heather, Mark Stehr, and Justin Sydnor. "Incentives, commitments, and habit formation in exercise: evidence from a field experiment with workers at a fortune-500 company." *American Economic Journal: Applied Economics* 7.3 (2015): 51-84.

Strathman, A., et al. "Considerations of future consequences (CFC scale)." Measurement Instrument Database for the Social Science. doi 10 (2013).

Thaler, R. H., & Sunstein, C. R. . *Nudge: Improving Decisions about Health, Wealth, and Happiness*. Yale University Press. (2008)

Appendices

Structure of the experiment over three weeks

Second week

Notification about changes in instructions, visible for 16 hours (+ email reminder)

Third week

Notification about changes in instructions visible for 16 hours (+ email reminder)

Today: 12:00
Demo

Real starts at 8:00 p.m

First week
21 rounds of
8 hours each

Second week
21 rounds of
8 hours each

Third week
21 rounds of
8 hours each

8 hours   16hours   16 hours

Payment=one week at random

Instructions

Welcome! This experiment today will last about 30 minutes and you will receive 5 Euro for your participation. If you want, you can then participate in a second phase of the experiment that will last 21 days. Depending on the choices you will make during these 21 days, you can earn other euros.

Please read these instructions carefully. This first phase will take place in this virtual room (including the demo). At the end of this, the real experiment will begin, taking place on your device and lasting 21 days.

IMPORTANT: We remind you that your participation will remain anonymous to the other participants as well as to the experiments. You will receive an identification number, automatically assigned by the computer, and it will be used for payments.

DESCRIPTIONS OF THE GAME

In this experiment you will be given 350 ECU to invest in three different stocks. One ECU corresponds to 0.04 Euro (that is 50 ECU= 2 Euro).

Your job is to choose when to buy and sell each stock, so that you earn the most money by the end of the experiment. Throughout the experiment, you will see the price of each stock changing (more detail below), and you will use this information to decide when to buy and sell. When you sell a stock, you receive an amount of cash equal to the price of the stock. When you buy a stock, you receive one unit of the stock, but you must give up an amount of cash equal to the current price of the stock.

The three stocks you can buy or sell are simply called Stock A, Stock B, and Stock C. At the beginning, the experiment each one of three stocks will be automatically assigned to you and each one costs 100 ECU. Therefore, at the beginning of the experiment you will have the following situation:

Stock	Quantity	Current Price	ECU Value	Euro Value
A	1	100	100	4
B	1	100	100	4
C	1	100	100	4
Cash		50	50	2
Total value			350	14

For the entire duration of the experiment, you can hold one unit at most of each stock. You cannot hold a negative quantity (that is you cannot sell stocks that are not at your disposal). Nevertheless, you might have a negative amount of cash. That will happen should you buy a stock at a price that is higher than the amount of cash you have at the moment of the purchase. This negative amount will be detracted from your earnings at the end of the experiment.

Structure of the market

In this experiment, <u>every day</u> you will be able to buy and sell stocks in three different time window that will call "sessions of the market":

1. **Morning session: from 4:00 a.m. to 12:00 p.m;**
2. **Afternoon session: from 12:00 p.m. to 8:00 p.m;**
3. **Night session: from 8:00 p.m. to 4:00 a.m.**

In particular, during each market sessions,

- the **price of each stock will be updated** and you will be informed whether the price increased or decreased, and by which amount;
- at the new price you will have the possibility to sell each stock (should you hold it) or buy it (should you not).

You will be able to make your choice at any moment during the opening of the market sessions but you cannot make a choice once the session is closed.

Structure of the game during 21 days

The game will be equally repeated with each three market sessions over 21 days, with little variations that will be introduced after 7 and 14 days that will be notified directly to your screen as well as by email (see further below "Earnings").

In particular, at the beginning of the second and third week you will receive a notification of the changes that will intervene during each week. This notification will remain visible on your screen for at least <u>16 hours</u>. Once this time expires, you will be able to play again. You will receive a reminder to your email as well.

How the stock price changes

Each stock changes price according to the exact same rule. Each stock is either in a good state or in a bad state. In the good state, the stock goes up with 70% chance, and it goes down with 30% chance. In the bad state, the stock goes down with 70% chance and it goes up with 30% chance.

Once it is determined whether the price will go up or down, the *size* of the change is always random, and will either be ECU 5, ECU 10, or ECU 15. For example, in the bad state, the stock will go down with 70% chance, and the amount it goes down by ECU 5, ECU 10, or ECU 15 with equal chance. Similarly, the good stock will go up with 70% chance, and the amount it goes up by will either be ECU 5, ECU 10, or ECU15.

The stocks will all randomly start in either the good state or bad state, and after each price update, there is a 20% chance the stock switches state.

The tables below summarise these information

Price changes

	Good state	*Bad state*
+ (UP)	70%	30%
- (DOWN)	30%	70%

State changes

	Good state today	*Bad state today*
Good state tomorrow	80%	20%
Bad state tomorrow	20%	80%

Earnings

You will play this game for 21 days in total, divided into three phases of 7 days each. In particular, at the beginning of each new phase (i.e. after 7 and 14 days) you will be able to again buy each stock at 100 ECU and the state of each stock will be restarted, i.e. randomly drawn again as at the beginning of the experiment.

Your earnings will be restarted as well at the beginning of each new phase (i.e. after 7 and 14 days) and will be computed for each phase at the end of the experiment. More precisely, the earnings corresponding to each phase will be equal to the amount of cash you accrued over the two scanning sessions from buying and selling stocks, plus the current price of any stocks that you own.

Earnings=cash + price A*(Hold A) + Price B*(Hold B) + Price C*(Hold C)

Finally, one phase of 7 days out of the three will be randomly selected for payment (i.e. you will be paid according to the total earnings of a randomly selected week).

Your final earnings will be converted in Euro at the exchange rate of 1 ECU= 0.04 Euro.

For payment you will have two options:

1. by IBAN
2. by cash at the Department of Economics (but only if compatible with the actual normative)

In any case, you will have to send an email to caterina.giannetti@gmail.com

The authors

Alexia Gaudeul

Research Associate

University of Goettingen, Germany

E-mail: alexia.gaudeul@wiwi.uni-goettingen.de

Caterina Giannetti,

Assistant Professor

University of Pisa, Italy

E-mail: caterina.giannetti@unipi.it