

Target Return and Negative Interest Rates*

C. D'Hondt[†], R. De Winne[‡], and A. Todorovic[§]

April 7, 2021

Abstract

In this paper, we question whether it is the zero threshold or the target return that impacts the most the risk-taking behavior of individuals when experiencing negative interest rates (NIRs). In our experiment, we attribute either low or high target returns to participants and ask them to make independent investment decisions when the risk-free rate fluctuates around their target returns, and for some of them, turns out negative. We find that the prevailing reference point is the target return, whatever its level. This result still holds even when the risk-free rate becomes negative, suggesting that the target return is a better driver of risk-taking than the zero-lower bound.

Keywords: *negative interest rates, risk-taking, target return*

JEL Classification: *G11, G21, G40, G41, G51, G53*

1 Introduction

Negative interest rates (NIRs hereafter) are no longer hypothetical. Since the 2008 financial crisis, NIRs have come true on some bond and monetary markets, and have even been set as policy rates by several central banks of industrialized countries (Altavilla et al., 2019; Brown, 2018).¹ To a lesser extent, a few commercial banks in Europe have also implemented NIRs and require some of their depositors to pay to let money on their savings accounts.² Unlike positive interest rates (PIRs hereafter) that deliver some gratification for being patient and postponing consumption, NIRs entail a kind of punishment aiming at stimulating economic growth through higher risk-taking and/or spending (Agarwal and Kimball,

*The authors gratefully acknowledge financial support from the Belgian French-speaking community (ARC grant 18/23-089). Any errors are the full responsibility of the authors.

[†]Louvain Finance (LIDAM), UCLouvain ; e-mail address: catherine.dhondt@uclouvain.be

[‡]Louvain Finance (LIDAM), UCLouvain ; e-mail address: rudy.dewinne@uclouvain.be

[§]Louvain Finance (LIDAM), UCLouvain ; e-mail address: Aleksandar.todorovic@uclouvain.be; Corresponding author

¹Until now, seven central banks have implemented this unconventional monetary policy measure: Bulgaria, Denmark, Euro Area, Hungary, Sweden, and Switzerland (Boungou, 2020).

²Some examples are the Alternative Bank Schweiz AG in Switzerland which has been one of the first banks to apply a deposit rate of -0.75%, the Danish Yske Bank that applies a deposit rate of -0.6%, the private bank Puilaetco Dewaay in Belgium which plans to apply NIRs for wealthy depositors, and the Dutch online broker DeGiro which wants to apply NIRs on deposits exceeding 2,500 EUR. According to Altavilla et al. (2019), pioneers in the implementations of NIRs are investment grade banks and more healthy banks, which start requiring their corporate depositors to pay for savings.

2015; Kimball, 2015). Because banks, firms, institutional investors or regular people are punished when hoarding money with NIRs, they should rather spend it, invest it, or lend it.

Low interest rates are typically associated with increased risk-taking among institutional investors, who "reach for yield" (Choi and Kronlund, 2018). This relationship has been confirmed when interest rates turn out negative. For example, institutional investors in money market funds must take more risks because NIRs prevent them from achieving positive returns through safe investments (Maggio and Kacperczyk, 2017). The impact of NIRs on banks' risk-taking is however more puzzling, since the banks in countries whose the monetary authority implemented NIRs have decreased their risk-taking, instead of having increased it (Boungou, 2020).³

What about the "reach for yield" among regular people? Past research is scarce but provides supportive evidence for it. When analyzing fund flows of retail investors between mutual funds, Leung and Zhu (2018) show that retail investors channel their funds according to the performance of mutual fund managers when interest rates are assumed to be normal (i.e., not zero nor negative). Nevertheless, as the interest rate becomes extraordinarily low, retail investors channel their funds towards the mutual funds with the risky benchmark. In the same vein, Ganzach and Wohl (2018) examine the impact of a decrease of the risk-free rate on individual investments between a risk-free asset and a risky asset. While keeping the risk premium constant, they find that decreasing the risk-free rate leads to an increase in risk-taking (i.e., more individuals opt for a risky asset when the interest rate on the risk-free asset decreases). Lian et al. (2018) obtain similar findings, and report that low interest rates are linked to higher money allocations to stocks and lower allocations to cash.

The specific case of zero or negative rates and their impact on the appetite of individuals for risk has however led to mixed results. Bracha (2020) compares investment decisions with both PIRs and NIRs and finds no significant difference for risk-taking. By contrast, Baars et al. (2020) indicate that the sign of interest rates plays a key role in risk-taking, i.e., a reduction in interest rates does not affect risk-taking in general; risk-taking increases significantly only if the interest rate drops below zero. More recently, David-Pur et al. (2020) report two key findings from a series of lab experiments analyzing both borrowing and investment behavior under PIRs and NIRs. First, the zero-interest rate has the strongest impact on individuals' investment decisions. Based on this result, David-Pur et al. (2020) conclude that zero interest rates are more efficient than NIRs to encourage individuals to either borrow money or take risks. Second, they find no statistical difference between the effect that PIRs and NIRs have on the change in the allocation of risky assets in investment portfolios, although the level of the allocation to risky assets does increase when interest rates move downwards.

In the aforementioned literature, the zero-level acting as a seemingly natural reference point might explain increased risk-taking with NIRs. Based on Prospect Theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) that considers the current level of wealth as reference point, individuals increase their risk-taking when the risk-free rate becomes negative to maintain their current level of

³Using a Diff-in-Diff approach to compare the risk-taking of banks whose monetary authorities implemented NIRs with the one of banks whose monetary authorities did not, Boungou (2020) finds that the risk-taking of the latter is higher. Moreover, among the banks whose monetary authorities implemented NIRs, both small and better-capitalized banks have taken less risk, compared to the large and less capitalized ones.

wealth and avoid losses. The concept of diminishing sensitivity entailed in Prospect Theory might also explain why the impact of a change is starker near the zero threshold than far from it (Ganzach and Wohl, 2018). Accordingly, a risk premium of 4% on a stock is more meaningful when the risk-free rate is 0% than when the risk-free rate is 10%.

In this paper, we contribute to this burgeoning literature on what stimulates the appetite of individuals for risk when experiencing NIRs. The novelty of our approach is to introduce a variable of interest in investment decisions, that is the target return. According to Payne et al. (1980), the latter can be defined as “the pre-determined benchmark return used by a decision-maker to translate monetary outcomes into gains and losses”. Because the target return represents an expectation of the return an investor wants to achieve, it is likely to impact the risk-taking behavior (Shefrin and Statman, 2000). Unlike the above papers that implicitly consider the zero threshold as the reference point by default (Bracha, 2020; Ganzach and Wohl, 2018; Lian et al., 2018; Baars et al., 2020; David-Pur et al., 2020), we run an experiment wherein explicit target returns are set and randomly attributed to participants. First, using Prospect Theory, we are able to check whether the target return serves as a reference point and effectively impacts the risk-taking behavior of individuals when making investment decisions. Next, and more importantly regarding NIRs, we determine which of the target return or the zero threshold is the prevailing reference point when the risk-free rate drops below zero.

Broadly speaking, we can relate target returns to mere goals, which are defined as a specific level of performance (Heath et al., 1999; Rechenberg et al., 2016).⁴ This paper builds on both Heath et al. (1999) and Larrick et al. (2009), who postulate that goals incorporate the three main principles of Prospect Theory, namely, reference point, loss aversion, and diminishing sensitivity. When target returns are explicitly set, we posit they act as goals that alter the psychological value of outcomes and impact the appetite for risk. To check this conjecture within the context of investment decisions, we derive several hypotheses based on goal-setting effects that predict changes in risk preferences, and we test them in a bespoke experiment.

Our experiment has a mixed design, wherein we manipulate the target return to check its impact on risk-taking. Specifically, we use two conditions (high or low target return). Participants are randomly assigned to one of these conditions and are subsequently requested to make five independent investment decisions. Two financial assets are available for all participants: a risk-free asset (i.e., a deposit account which delivers the risk-free rate) and a risky asset (i.e., a stock index fund with a constant risk premium - 5% above the risk-free rate). We manipulate the risk-free rate to have it both above and below target returns. These variations in the risk-free rate are similar in both conditions. However, participants in the High-Target-Return condition face risk-free rates that always remain positive, while participants in the Low-Target-Return condition experience both positive and negative risk-free rates. Participants are informed that the fluctuations in the interest rates on the deposit account depend on the monetary policy of the Central Bank, while the stock index fund returns vary with the current market conditions. After

⁴Past research has found that mere goals have an impact on both performance and risk-taking (Locke and Latham, 1990). In particular, challenging goals have a larger positive impact on performance than less challenging goals, even when goals are extremely high (Heath et al., 1999). Specific and challenging goals also increase performance more than “do your best” goals.

having made their investment decisions, participants are asked to complete a short questionnaire. The latter allows us to collect individual information about their target return in reality, their subjective risk-taking inclination, as well as whether they own(ed) a savings account and an actual portfolio of financial assets.

Our original approach, which pioneers in using target returns in both PIR and NIR contexts, provides insightful results about the “reach for yield” among regular people. First, we show that explicit target returns significantly drive risk-taking in investment decisions. Building on Heath et al. (1999) and Larrick et al. (2009), our findings confirm that the target return is the reference point of participants when making investment decisions. This holds whatever the level of the target return, i.e., in both High- and Low-Target-Return conditions. Participants invest effectively more in the stock index fund when the risk-free interest rate is below their target return than when it is above their target return. Consistent with loss aversion, participants’ reactions when the risk-free rate falls below their target return (i.e., increase in risk-taking) exceed the corresponding reactions when the risk-free rate goes above their target return. Regarding diminishing sensitivity that predicts a greater impact of a given shift in the risk-free rate when occurring near the target return (than far from it), we only find partial evidence. Participants exhibit the expected behavior only when the risk-free rate fluctuates below their target return.

Of particular interest is our result for the specific case when the risk-free rate drops below zero in the Low-Target-Return condition. Comparing the average increases in risk-taking for similar shifts of the risk-free rate in both conditions (i.e., from 0% to -2% in the Low-Target-Return condition and from 4% to 2% in the High-Target-Return condition), we find no significant difference. This means that the target return still acts as the prevailing reference point for our participants who face a negative risk-free rate. In other words, having the risk-free rate crossing the zero-level does not especially affect risk-taking when people have an explicit target return (that differs from zero); people continue to adjust their allocation to the risky asset based on the distance to their target return. The larger the distance between the target return and the risk-free rate, the lower is the marginal increase in risk-taking (because of diminishing sensitivity). We conclude that the move of the risk-free rate into negative territory does not activate any additional (i.e., extra) risk-taking.

When further investigating the impact of the target return on risk-taking using panel data regression models, the expected risk-seeking behaviors are confirmed when the risk-free rate is below the target return. Whatever the condition (i.e., the level of target return), when the risk-free rate is 2% (4%) below the target return, the allocation to the risky asset decreases by about 20% (23%). Such reductions are both statistically and economically significant. Furthermore, when adding an interaction variable that refers to situations when participants in the Low-Target-Return condition face a negative risk-free rate, our results confirm that the target return still plays the crucial role of reference point and drives risk-taking.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature and presents our hypotheses. Section 3 describes our experimental setting. Section 4 reports the results. Section 5 concludes.

2 Literature and Hypotheses

2.1 Prospect Theory and Cumulative Prospect Theory

To formalize the impact of target returns on the risk-taking behavior, we build on Prospect Theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). This well-known model of decision-making under risk entails a value function derived from the changes in wealth, rather than ultimate wealth.⁵ Prospect Theory is often summarized in its three core principles : reference point, loss aversion, and diminishing sensitivity.

The crucial notion of reference point implies that the value of any outcome is a function of both the current position (i.e., the reference point) and the shift (either positive or negative) from it. In short, this value function is concave for gains (positive shifts) and convex for losses (negative shifts). In addition, outcomes are multiplied with decision weights, which refer to the transformation of probabilities in Cumulative Prospect Theory. More precisely, the probability weighting function is a rising reverse S-shaped function, such that small (high) probabilities are over- (underweighted), while the ones in the middle match decision weights.⁶

Loss aversion is related to the certainty effect, i.e., people overweigh outcomes with certainty. This leads to risk aversion in choices with sure gains and risk-seeking in choices with sure losses. In other words, people prefer sure gains over probable gains in the positive region, while they prefer probable losses over sure losses in the negative region.⁷ Loss aversion is usually encapsulated in the expression “losses loom larger than gains”, i.e., the negative subjective value given to a loss is higher than the positive subjective value given to an equivalent gain.⁸ Using the shapes of both the value function and the probability weighting function, Tversky and Kahneman (1992) highlight the following key patterns : risk aversion for gains of high probability (i.e., sure gains); risk-seeking for losses of high probability (i.e., sure losses); risk-seeking for gains of low probability; risk aversion for losses of low probability.

The third core principle of Prospect Theory is diminishing sensitivity, i.e., the marginal value of both gains and losses decreases with their magnitude. Consistent with a concave value function, the difference between an increase from 100 to 200 is greater than an increase from 1,000 to 1,100. Similarly, the difference between a loss of 100 to 200 is greater than the equivalent increase in loss from 1,000 to 1,100, which is consistent with a convex value function.

2.2 Goals as reference points

Heath et al. (1999) postulate that goals serve as reference points and change the value of outcomes as Prospect Theory predicts. Specifically, these authors focus on “mere” goals, which are defined as

⁵For comparison, the value function of the Expected Utility Theory is a function of the final assets and is a concave function of wealth.

⁶Although this probability weighting function is circumscribed independently for gains and losses, it exhibits similar curvature (see Figure 4 in Appendix A).

⁷For example, people prefer to gain 3.000 for sure over 4.000 with 0.8 probability of occurrence. But they prefer the loss of 4.000 with 0.8 probability of occurrence over the sure loss of 3.000.

⁸For example, the negative subjective value given to a loss of 100 seems two times higher than the positive subjective value given to an equivalent gain.

specific levels of performance with no discrete and discontinuous payoff.⁹ To check their predictions, Heath et al. (1999) use several short experiments wherein people who have pre-determined workout plans with different goals are depicted. When asking participants different questions aiming at describing the emotions of these people depicted in various situations, Heath et al. (1999) demonstrate not only that goals, like reference points, divide outcomes into regions of good and bad (i.e., success and failure), but also that the valuation of outcomes relative to a goal can be explained by loss aversion and diminishing sensitivity. Furthermore, Heath et al. (1999) show that the principles of the value function are sufficient to explain a wide range of empirical findings in the goal-setting literature on motivation.

Building on Heath et al. (1999), Larrick et al. (2009) focus on how goals induce risk-taking in two empirical contexts, namely, negotiation and decision-making. In particular, Larrick et al. (2009) insist on the fact that a specific goal serves as reference point, creating a region of perceived losses for outcomes below that goal. To check whether specific goals increase risk-taking, these authors run three experiments wherein they compare a "your best" condition to a "specific, challenging goal" condition.¹⁰ Their first two experiments show that goal-setting leads to more risk-taking in simple distributive bargaining tasks. They find consistent results in their third experiment wherein the impact of goal-setting is tested in a one-shot gambling task with clearly stated probabilities and payoffs. Larrick et al. (2009) conclude that all their findings point to increased risk-taking with specific goals, supporting their predictions based on Prospect Theory.

In the present paper, we consider a typical goal in the context of investment decisions, that is the target return. The latter can be defined as an outcome towards which an individual strives while making investment decisions among risky choices. According to Payne et al. (1980), such a goal provides any investor with the opportunity to perceive the outcome of his/her investment as a gain or a loss. Hence, all outcomes above (below) the target should be perceived as gains (losses). When target returns are explicitly set, we posit they act as specific goals that alter the psychological value of outcomes and impact the appetite for risk. To check this conjecture, we formulate in Section 2.3 several hypotheses based on goal-setting effects and expected changes in risk preferences.

2.3 Hypotheses

To facilitate the understanding of our hypotheses, we use Figure 1 that illustrates three different value functions for outcomes about the risk-free rate based on Prospect Theory. In the first (green) value function, the target return is set to 0% (i.e., zero-level). In the second (blue) value function, the target return is set 2% (i.e., at the low level) while in the third (red) value function, the target return is set at 6% (i.e., at the high level). In accordance with the three core principles described in Section 2.1, all these three value functions display an S-shape, i.e., they are concave when risk-free rates exceed the target return and convex when risk-free rates fall short of the target return. In addition, when risk-free rates are below target returns, all the three value functions are steeper, compared to when risk-free rates are

⁹According to Heath et al. (1999), risk-taking in situations wherein people have goals tied as discrete and discontinuous external rewards (e.g., promotion or bonus), can be easily explained by basic economic calculations. By contrast, increased risk-taking in response to mere goals is a fundamentally psychological phenomenon.

¹⁰Larrick et al. (2009) present these conditions as the standard manipulation in the goal-setting literature.

above target returns. In absolute terms, negative subjective values when risk-free rates are below target returns exceed the positive subjective values when risk-free rates exceed target returns for corresponding magnitudes. Next, the steepness of these value functions is high (low) when risk-free rates are near (far from) target returns, i.e., the impact of a change in the risk-free rate diminishes when moving further from the target return. In other words, an increase in the risk-free rate above the target return by a given magnitude (for example, 2%) has a stronger impact when it happens near the target return than when it is distant from it. Similarly, a larger negative impact of a decrease in the risk-free rate is assumed when occurring near the target return (compared to when it happen far from it).




Figure 1: Zero-level, Low-Target-Return, and High-Target-Return value function for risk-free rates

This figure depicts three value functions for outcomes about the risk-free rate based on Prospect Theory. In the first (green) value function, the target return is set to 0% (i.e., zero-level). In the second (blue) value function, the target return is set 2% (i.e., at the low level) while in the third (red) value function, the target return is set at 6% (i.e., at the high level). The horizontal axis refers to the risk-free interest rate, while the vertical axis refers to its subjective values based on Prospect Theory.

Building on the primary role of the reference point (i.e., flagging outcomes as gains or losses), our first hypothesis postulates risk aversion when the risk-free rate exceeds the target return and risk-seeking when the risk-free rate drops below the target return. In Figure 1, when the risk-free rate exceeds 2% in the Low-Target-Return condition (points D and E) and 6% in the High-Target-Return condition (points I and J), a low proportion of money invested in the risky asset is expected. On the contrary, when the

risk-free rate falls short of 2% in the Low-Target-Return condition (points B and C) and 6% in the High-Target-Return condition (points G and H), participants are expected to be risk-seeking, which should be reflected in a higher proportion of money invested in the risky asset.

Our second hypothesis is driven by loss aversion, i.e., losses loom larger than gains. When the risk-free rate is below the target return, the subjective loss associated with this outcome should exceed the subjective gain when the risk-free rate is above the target return by the same magnitude. Therefore, we expect that the increase in the proportion invested in the risky asset when the risk-free rate implies a given loss (e.g., -2%) will be higher than the decrease in the proportion invested in the risky asset when the risk-free rate implies a gain of similar magnitude (e.g., +2%). Put differently, the higher unattractiveness of the risk-free asset when entailing a loss (e.g., point B or C in the Low-Target-Return condition and point G or H High-Target-Return condition) should lead to a larger increase in the amount invested in the risky asset than the corresponding decrease of the amount invested in the risky asset when the risk-free asset implies a gain (e.g., point D or E in the Low-Target-Return condition and point I or J High-Target-Return condition).

Our third hypothesis is based on diminishing sensitivity, i.e., the impact of a given change in gains and losses diminishes as we move further from the reference point. Accordingly, we hypothesize that the impact of a shift in the risk-free rate near the target return will lead to a higher change in risk-taking, compared to any equivalent shift in the risk-free rate occurring far from the target return. When the risk-free rate is just above the target return (i.e., when the risk-free rate moves from 2% to 4% (from point A to D) in the Low-Target-Return condition, and from 6% to 8% (from point F to I) in the High-Target-Return condition), the decrease in the proportion of money invested in the risky asset should be significantly higher than the decrease in the proportion of money invested in the risky asset when the risk-free rate is far above the target return (from 4% to 6% (from point D to E) in the Low-Target-Return condition, and 8% to 10% (from point I to J) in the High-Target-Return condition). Symmetrically, when the risk-free rate drops just below the target return, the increase in the proportion of money invested in the risky asset (i.e., when the risk-free rate moves from 2% to 0% (from point A to B) in the Low-Target-Return condition, and from 6% to 4% (from point F to G) in the High-Target-Return condition) should be significantly higher than the increase in the proportion of money invested in the risky asset when the risk-free rate drops far below the target return (from 0% to -2% (from point B to C) in the Low-Target-Return condition, and from 4% to 2% (from point G to H) in the High-Target-Return condition).

Our last hypothesis posits that the target return is the prevailing reference point, even when the risk-free rate turns out negative. Looking at Figure 1, we need to compare the change in the proportion of money invested in the risky asset when the risk-free rate drops from 0% to -2% (from point B to C) in the Low-Target-Return condition, with the change in the proportion of money invested in the risky asset when the risk-free rate drops from 4% to 2% (from point G to H) in the High-Target-Return condition. In both situations, the risk-free rate is below the target return and the distance between the risk-free rate and the target return is the same. If the target return still prevails as reference point when the risk-free rate drops below zero, participants in both target return conditions should exhibit similar risk-seeking behavior, i.e., the amount invested in the risky asset should increase in the same proportion.

3 Experimental setting

We use oTree (Chen et al., 2016) to design our experiment, with one between and two within factors. The between factor is the target return condition to which each participant is randomly assigned. The target return is set to 6% in the High-Target-Return condition and to 2% in the Low-Target-Return condition (see Table 1). Whatever its level, this target return is kept constant for the whole experiment. All participants are informed that they are going to make several independent investment decisions, with two alternatives available: a deposit account and a stock index fund. The riskiness of both options is explicitly explained, i.e., the deposit account is presented as the risk-free asset and the stock index fund as the risky asset. The two within factors are therefore the interest rate on the deposit account (i.e., risk-free rate) and the stock index fund return. Participants are instructed that (1) the interest rate on the deposit account can be either positive or negative, depending on the monetary policy set by the Central Bank, and (2) the stock index fund return fluctuates depending on the market conditions. The currency in this experiment is the ECU (Experimental Currency Unit), and the exchange rate with the US dollar (\$) is 1:1.

In total, each participant is asked to make five independent investment decisions. For each of them, participants in the Low-Target-Return (High-Target-Return) condition have to allocate an identical amount of 50 (30) ECU between the two available assets.¹¹ The total amount invested in each decision must add up to 100%, i.e., they need to invest all their money. After each investment decision, participants learn about their performance, i.e., they observe their realized returns on both the deposit account and the stock index fund. More specifically, participants are provided with information on the global performance for both assets in both absolute and relative terms, as well as on the performance for each asset separately.

Table 1 shows how the risk-free rate fluctuates across investment decisions. In each condition, two risk-free rates exceed the target return (see the first two rows in each panel), one risk-free rate just matches the target return (see the third row in each panel), and two risk-free rates fall short of the target return (see the last two rows in each panel). In the High-Target-Return condition, the risk-free rate is set to 10%, 8%, 6%, 4% and 2%, while it is set to 6%, 4%, 2%, 0% and -2% in the Low-Target-Return condition. It is noteworthy that, in both conditions, each of these values is randomly assigned to a given investment decision.¹² Depending on the risk-free rate, the expected return of the stock index fund also varies across investment decisions, with a constant risk premium at 5% (see the penultimate column in Table 1). For example, when the risk-free rate decreases by 2%, the stock index fund expected return goes down by 2% as well. The realized return of the stock index fund within each investment decision takes one of the five possible outcomes displayed on Figure 3. Specifically, two outcomes exceed the expected return (i.e., 20% and 25%), one outcome just matches it (i.e., 15%), and two outcomes fall short of it

¹¹The amounts to invest vary across the two conditions to keep the final payoffs of participants similar in both conditions. Since the interest rates on the deposit account and the expected returns on the stock index fund are by default higher in the High-Target-Return condition, we reduce the initial amount in this condition. We follow this procedure from Bracha (2020) and Ganzach and Wohl (2018), who also define the available amounts to equalize the final payoffs in their experiments.

¹²This means that the risk-free rate is not monotonously increasing or decreasing across the five investment decisions.

(i.e, 5% and 10%). The standard deviation of the distribution of the stock index fund returns is kept constant at 6.3% across investment decisions in both conditions (see the last column in Table 1).

Table 1: Investment alternatives across decisions

Panel A: High-Target-Return condition						
Decision	TR	R_F	R_F-TR	$E(R_s)$	RP_s	STD_s
1	6%	10%	4%	15%	5%	6.3%
2	6%	8%	2%	13%	5%	6.3%
3	6%	6%	0%	11%	5%	6.3%
4	6%	4%	-2%	9%	5%	6.3%
5	6%	2%	-4%	7%	5%	6.3%
Panel B: Low-Target-Return condition						
Decision	TR	R_F	R_F-TR	$E(R_s)$	RP_s	STD_s
1	2%	6%	4%	11%	5%	6.3%
2	2%	4%	2%	9%	5%	6.3%
3	2%	2%	0%	7%	5%	6.3%
4	2%	0%	-2%	5%	5%	6.3%
5	2%	-2%	-4%	3%	5%	6.3%

This table describes the investment alternatives across decisions in High- and Low-Target-Return conditions, respectively. TR refers to the target return, R_F to the interest rate on the deposit account (i.e., risk-free rate), R_F-TR to the difference between the risk-free rate and the target return, $E(R_s)$ to the expected return on the stock index fund, RP_s to the risk premium of the stock index fund, and STD_s to the standard deviation of the stock index fund returns.

To help participants understand the investment alternatives available at each decision, we provide them with graphical representations similar to Figures 2 and 3. Figure 2 shows the interest rate on the deposit account as well as the probability of achieving it (that is always equal to 100%). Figure 3 presents the possible returns on the stock index fund with their respective probabilities. These figures should help participants directly compare possible outcomes and probabilities of occurrence.¹³

¹³The advantage of such graphical representations is a better understanding for participants as well as a decrease in uncertainty regarding the available alternatives. However, we are aware that these graphical representations might induce some potential framing effect, probably in favor of the risk-free asset.




Figure 2: Interest rate on the deposit account

This figure shows the interest rate on the deposit account in a randomly chosen investment decision. The horizontal axis refers to the interest rate on the deposit account, while the vertical axis refers to the probability of occurrence for this risk-free rate (which is 100% by definition).




Figure 3: Distribution of the stock index fund returns

This figure presents the distribution of the stock index fund returns in a randomly chosen investment decision. The horizontal axis refers to the five possible outcomes for the stock index fund, while the vertical axis refers to their respective probabilities of occurrence.

Our experiment is incentivized, with a reward system made of a fixed amount of \$1.50 plus a premium. The former is received by each participant, whatever the results of his/her decisions. By contrast, the premium is an additional amount of money that directly depends on his/her average performance across the five investment decisions. The amount of this premium is simply equal to the average realized return multiplied by \$50 in the Low-Target-Return condition or \$30 in the High-Target-Return condition. For example, if a participant in the Low-Target-Return condition gets an average realized return equal to 2%, his/her reward would be \$2.5, which consists of a premium of \$1 plus the fixed amount (\$1.5). The same average realized return in the High-Target-Return condition leads to a premium of \$0.6 ($\$30 \times 2\%$) and a total reward of \$2.1.

After the decision-making part, participants are requested to answer a hypothetical question about what would be their own target returns in reality. Specifically, we ask them to declare the target return they would set if they had an actual portfolio of financial assets. The motivation behind this question is to check for any anchoring towards the target return of 6% and 2% in High- and Low-Target-Return conditions, respectively. To end up, participants have to fill in a small questionnaire made of three parts. The first part elicits their subjective risk-taking inclination with the DOSPERT scale (Blais and Weber, 2006), restricted to the six questions about gambling and investing (details are available in Appendix B). The second part of the questionnaire investigates the actual savings behavior of participants, i.e., whether they have or had in the past money on their savings account. The last question focuses on whether participants hold or held in the past an actual portfolio of financial assets. Our main purpose is

to check whether such individual information help explain heterogeneity in behaviors across participants.

4 Results

4.1 Descriptive statistics

We recruited 91 participants¹⁴ via Prolific Academic, which is a crowdsourcing platform for online research. Using the available prescreening options, participants are requested to be US citizens whose first language is English, and who are between 18 and 65 years old.

Table 2 provides descriptive statistics for the full sample as well as for the High-Target-Return and Low-Target-Return conditions, respectively. Focusing on the full sample, the average age of our participants is 32 (with a median of 30). The sample is well-balanced with 51% of males (and 49% of females). One quarter of the participants are students, while 74% of them are employed. On average, our participants report they would set their target return at 15% in reality. The median is however much smaller at 6%. The average risk tolerance across all participants is about 2.76 on a scale ranging from 0 (extremely unlikely to take a risk) to 7 (extremely likely to take a risk).¹⁵ This indicates that the typical participant is rather risk-averse. Interestingly, 90% of the participants own an actual deposit account, while 55% of them also own an actual portfolio of financial assets.

Since participants are randomly assigned to the High- or Low-Target-Return condition, the two subsamples are overall similar in terms of age, gender, and status. We note however that participants in the High-Target-Return condition are on average slightly more likely to be employed (77%), compared to participants in the Low-Target-Return condition (72%). As expected, our manipulation leads to some anchoring since participants in the High-Target-Return condition declare on average a higher target return in reality (20% versus 11%).¹⁶ As for risk tolerance, the average score is somewhat lower in the High-Target-Return condition (2.67 versus 2.86) while medians are just identical (at 2.83). On average, 86% (95%) of participants in the High- (Low-Target-Return) condition own a real deposit account. More than half of participants in both conditions (53% in the High- and 56% in the Low-Target-Return condition) invest in real financial assets.

¹⁴We invited 100 persons to participate in our experiment, but 9 participants dropped off after reading the instructions.

¹⁵Based on the answers to the six questions taken from the DOSPRT scale (available in Appendix B), we calculate risk tolerance as the average of the scores on these questions.

¹⁶This difference is statistically significant at the 1% level (see Appendix C).

Table 2: Descriptive statistics about the sample of participants

Variable	Full sample				High-Target-Return condition				Low-Target-Return condition			
	Mean	Median	Q1	Q3	Mean	Median	Q1	Q3	Mean	Median	Q1	Q3
Age	32	30	25	37	32	29	24	37	31	30	25	36
Male (in %)	51				52				50			
Student (in %)	25				28				22			
Employee (in %)	74				77				72			
Target return in reality (%)	15	6	5	10	20	6	5	20	11	5	3	9
Risk tolerance (0-7 scale)	2.76	2.83	2.17	3.33	2.67	2.83	2.00	3.33	2.86	2.83	2.17	3.50
Deposit account owner (in %)	90				86				95			
Owner of financial assets (in %)	55				53				56			

This table reports cross-sectional descriptive statistics (mean, median, lower and upper quartiles) about the participants in the sample. These statistics are computed both across the full sample as well as across the High-Target-Return and Low-Target-Return conditions separately. 'Target return in reality' refers to the target return participants would set if they had an actual portfolio of financial assets. 'Risk tolerance' refers to the subjective risk-taking inclination elicited using the six questions about gambling and investing in the DOSPERT scale. This measure of risk tolerance is calculated as the average of the scores on these questions, on a scale ranging from 0 (extremely unlikely to take a risk) to 7 (extremely likely to take a risk). 'Deposit account owner' is a binary variable set to one for participants who have or had in the past money on their savings account. 'Owner of financial assets' is another binary variable set to one for participants who hold or held in the past an actual portfolio of financial assets.

Across all decisions and both conditions, the average proportion of money allocated to the stock index fund is about 70% while the corresponding median reaches 83%. As expected, participants in the Low-Target-Return condition invest on average more in the risky asset (76%), compared to the participants in the High-Target-Return condition (64%). Table 3 provides further information about the average proportions of money allocated to the risky asset depending on the risk-free rate in each condition. In the High-Target-Return condition, the average proportion invested in the risky asset ranges from 42.39% when the risk-free rate is equal to 6% to 63.62% when the risk-free rate is set at 2%. The range is somewhat higher in the Low-Target-Return condition where the highest average proportion allocated to the risky asset is 76% when the risk-free rate is negative (-2%) while the lowest average proportion is 46.84% when the risk-free rate is set at 6%. In both conditions, the average allocation to the risky asset does not increase monotonically as the risk-free rate decreases, though an upward trend is noticeable.

Table 3: Allocation to the risky asset (in %)

Condition	Risk-free rate						
	-2%	0%	2%	4%	6%	8%	10%
High-Target-Return			63.62	59.28	42.39	49.35	50.51
Low-Target-Return	76.00	70.00	48.56	50.65	46.84		

This table presents the proportion of money allocated to the risky asset (i.e., the stock index fund with a constant risk premium of 5%) depending on the risk-free rate in the High-Target-Return condition and Low-Target-Return condition, respectively.

4.2 Univariate testing of our hypotheses

Our first hypothesis postulates risk aversion when the risk-free rate exceeds the target return and risk-seeking when the risk-free rate drops below the target return. Accordingly, participants in both conditions are expected to invest significantly more in the risky asset when the risk-free rate falls short of their target return than when the risk-free rate exceeds it. To test this hypothesis, we conduct one-sided paired t-tests for mean comparison in the proportions of money allocated the risky asset when the risk-free rate is both 2% and 4% below and above the target return. We report the results in Table 4, in Panel A when the risk-free rate is set at 2% below and above the target return and in Panel B when the risk-free rate is set at 4% below and above the target return. All the findings confirms our first hypothesis, i.e., participants in both conditions invest significantly more in the risky asset when the risk-free rate is below their target return than when it is above. When the risk-free rate is 2% below the target return, participants in the High- (Low-Target-Return) condition invest on average 9.93% (19.35%) more in the risky asset than when the risk-free rate is 2% above the target return. These differences are stronger when the risk-free rate is 4% below the target return, i.e., 13.11% (29.16%) in the High- (Low-Target-Return) condition.

Table 4: Risk-taking depending on whether the risk-free rate is above or below the target return

	High-Target-Return condition	Low-Target-Return condition
Panel A : 2% above and 2% below		
R_F at 2% below	59.28%	70.00%
R_F at 2% above	49.35%	50.65%
Difference	9.93%	19.35%
t-value	2.29	4.01
p-value	0.0134	<0.0001
Panel B : 4% above and 4% below		
R_F at 4% below	63.62%	76%
R_F at 4% above	50.51%	46.84%
Difference	13.11%	29.16%
t-value	1.93	4.81
p-value	0.0298	<0.0001

This table reports the average proportions of money invested in the risky asset depending on whether the risk-free rate (R_F) is above or below the target return, the differences between these average proportions, and the corresponding t-values and p-values from one-sided paired t-tests. Panel A refers to situations when the risk-free rate is 2% above and below the target return, while Panel B refers to situations when the risk-free rate is 4% above and below the target return.

To test our next hypotheses, we need to consider variations in risk-taking (instead of raw allocations to the risky asset). Therefore we compute all the variations observed in the proportions invested in the stock index fund for any possible shift in the risk-free rate in both conditions. We report the cross-sectional average variations in Table 5, in Panel A for the High-Target-Return condition and in Panel B for the Low-Target-Return condition. To make interpretations easier, Table 5 also provides the expected direction of these variations (i.e., increase or decrease in risk-taking) as well as the corresponding shift on Figure 1. Looking at this big picture, we can see that the expected increase in risk-taking when the risk-free rate moves below the target return occurs in both conditions. These average increases in risk-taking statistically differ from zero when considering shifts of both 2% and 4% in the risk-free rate near the target return. When the risk-free rate moves by 2% below but far from the target return, the average variations in the proportion of money allocated to the risky asset are still positive, though not statistically different from zero. These findings show that the decisions made by our participants when the risk-free rate is below the target return overall could fit with the value functions in the domain of losses depicted in Figure 1. By contrast, when the risk-free rate moves above the target return, we find average variations in risk-taking that never statistically differ from zero. Such results indicate that the decisions made by our participants when the risk-free rate is above the target return do not fit with the upper part of the value functions depicted in Figure 1. They are rather consistent with value functions that would be flat in the domain of gains.

Table 5: Cross-sectional average variations in risk-taking

Shift in R_F	Expected direction	Shift on Figure 1	Mean variation
Panel A: High-Target-Return condition			
$\Delta 2\%$ below and near	+	F $\rightarrow$ G	16.89%***
$\Delta 4\%$ below	+	F $\rightarrow$ H	21.23%***
$\Delta 2\%$ below and far	+	G $\rightarrow$ H	4.35%
$\Delta 2\%$ above and near	-	F $\rightarrow$ I	6.96%
$\Delta 4\%$ above	-	F $\rightarrow$ J	8.12%
$\Delta 2\%$ above and far	-	I $\rightarrow$ J	1.16%
Panel B: Low-Target-Return condition			
$\Delta 2\%$ below and near	+	A $\rightarrow$ B	21.44%***
$\Delta 4\%$ below	+	A $\rightarrow$ C	27.44%***
$\Delta 2\%$ below and far	+	B $\rightarrow$ C	6.00%
$\Delta 2\%$ above and near	-	A $\rightarrow$ D	2.09%
$\Delta 4\%$ above	-	A $\rightarrow$ E	-1.72%
$\Delta 2\%$ above and far	-	D $\rightarrow$ E	-3.81%

This table reports the cross-sectional average variations in the proportions invested in the stock index fund for any possible shift in the risk-free rate (R_F) in both conditions. Panel A refers to the High-Target-Return condition while Panel B refers to the Low-Target-Return condition. For each possible shift in the risk-free rate, the expected direction of these variations (i.e., increase or decrease in risk-taking) as well as the corresponding shift on Figure 1 are provided. ***, **, and * indicate whether these mean variations statistically differ from zero at 1%, 5%, and 10%, respectively.

Our second hypothesis is related to loss aversion. Looking at the value functions on Figure 1, their steepness below the target return (e.g., when the risk-free rate moves from point A to B or from A to C on the low target return value function) is higher than above the target return (when the risk-free rate moves from point A to D or from A to E on the low target return value function). When losses loom more than equivalent gains, participants' negative reaction when the risk-free rate moves below the target return is expected to outreach their positive reaction for any equivalent shift in the risk-free rate above the target return. In our particular setting, this means that the increase in the proportions invested in the risky asset when the risk-free rate drops below the target return (e.g., from point A to point B or C on the low target return value function) should exceed the decrease in the proportions allocated to the risky asset when an equivalent move of the risk-free rate occurs above the target return (e.g., from point A to point D or E on the low target return value function). This hypothesis can be tested in two ways: when the risk-free rate moves near the target return ($\pm 2\%$) and far from target return ($\pm 4\%$). Since our second hypothesis refers to the magnitude of the variations (and not the direction of them), we multiply the variations in the proportion invested in the risky asset above the target return by -1. Table 6 reports the results of one-sided paired t-tests for mean comparison in the variations of the allocation to the risky asset when the risk-free rate moves below and above the target return. In

both conditions, we find significant evidence of loss aversion. In Panel A, when the risk-free rate moves 2% below the target return, participants in the High-Target-Return condition increase on average their allocation to the risky asset by 16.89%, while they do not really change on average their allocation to the risky asset when the equivalent shift in the risk-free rate occurs above the target return. The difference in these average variations of risk-taking (23.85%) is highly significant. The findings are similar in the Low-Target-Return condition, i.e., a shift of the risk-free rate near the target return leads to a higher variation in the proportions of money allocated to the risky asset when it occurs below the target than when it happens above the target return (in the latter case, the average variation does not really differ from zero). We also find significant differences between the average variations in the allocation to the risky asset when the risk-free rate moves 4% below and above the target return (see Panel B). Our second hypothesis is therefore validated.

Table 6: Variations in risk-taking for loss aversion

	High-Target-Return condition	Low-Target-Return condition
Panel A : Shift near the target return ($\pm 2\%$)		
$\Delta 2\%$ below	16.89%***	21.44%***
$\Delta 2\%$ above	-6.96%	-2.09%
Difference	23.85%	23.54%
t-value	2.75	3.14
p-value	0.0043	0.0016
Panel B : Shift far from the target return ($\pm 4\%$)		
$\Delta 4\%$ below	21.23%***	27.44%***
$\Delta 4\%$ above	-8.12%	1.72%
Difference	29.35%	25.72%
t-value	3.05	3.40
p-value	0.0019	0.0007

This table reports the average variations in the allocation to the risky asset when the risk-free rate moves either near ($\pm 2\%$) or far from ($\pm 4\%$) the target return, the differences between these average variations, and the corresponding t-values and p-values from one-sided paired t-tests. Panel A refers to situations when the risk-free rate moves near the target return by 2%, while Panel B refers to situations when the risk-free rate moves far from the target return by 4%. Since we focus on the magnitude of the variations (and not the direction of them), we multiply the variations in the proportion invested in the risky asset above the target return by -1. *** indicates when the average variations statistically differ from zero (at the 1% level).

Our third hypothesis posits diminishing sensitivity. Looking at Figure 1, when the risk-free rate moves from 2% to 4% (from point A to D) in the Low-Target-Return condition, and from 6% to 8% (from point F to I) in the High-Target-Return condition, the decrease in the proportion of money invested in the risky asset should be significantly higher than the decrease in the proportion of money invested in the risky asset when the risk-free rate is far above the target return from 4% to 6% (from point D to E) in the Low-Target-Return condition, and 8% to 10% (from point I to J) in the High-Target-Return condition.

Symmetrically, when the risk-free rate drops just below the target return, the increase in the proportion of money invested in the risky asset (i.e., when the risk-free rate moves from 2% to 0% (from point A to B) in the Low-Target-Return condition, and from 6% to 4% (from point F to G) in the High-Target-Return condition) should be significantly higher than the increase in the proportion of money invested in the risky asset when the risk-free rate drops far below the target return (from 0% to -2% (from point B to C) in the Low-Target-Return condition, and from 4% to 2% (from point G to H) in the High-Target-Return condition). We find partial evidence of diminishing sensitivity in Table 7. Participants exhibit the expected behavior when the risk-free rate moves below the target return. In Panel A, when the risk-free rate drops just below the target return by 2% in the High-Target-Return condition (i.e., from 6% to 4%), participants increase on average their allocation to the risky asset by 16.89%. By contrast, the average marginal increase in risk-taking is equal to 4.35% when the risk-free rate moves from 4% to 2%. Similarly, when the risk-free rate drops just below the target return by 2% in the Low-Target-Return condition (i.e., from 2% to 0%), participants increase on average their allocation to the risky asset by 21.44%. In comparison, the average marginal increase in risk-taking is equal to 6.00% when the risk-free rate moves from 0% to -2%. In both conditions, these differences in the variations of risk-taking (12.54% and 15.44%) are statistically significant. However, consistent with Table 5 showing that the decisions made by our participants when the risk-free rate is above the target return do not fit with a concave value function in the domain of gains, the results reported in Panel B of Table 7 are not statistically significant. Consequently, our third hypothesis is only partially validated.

Table 7: Variations in risk-taking for diminishing sensitivity

	High-Target-Return condition	Low-Target-Return condition
Panel A : Equivalent shifts below the target return		
$\Delta 2\%$ near	16.89%***	21.44%***
$\Delta 2\%$ far	4.35%	6.00%
Difference	12.54%	15.44%
t-value	1.69	1.99
p-value	0.0485	0.0026
Panel B : Equivalent shifts above the target return		
$\Delta 2\%$ near	6.96%	2.09%
$\Delta 2\%$ far	1.16%	-3.81%
Difference	5.80%	5.90%
t-value	0.78	0.82
p-value	0.2187	0.2087

This table reports the average variations in the allocation to the risky asset when the risk-free rate moves by 2% near or far from the target return, the differences between these average variations, and the corresponding t-values and p-values from one-sided paired t-tests. Panel A refers to situations when the risk-free rate moves by 2% below the target return, while Panel B refers to situations when the risk-free rate moves by 2% above the target return. *** indicates when the average variations statistically differ from zero (at the 1% level).

Our last hypothesis assumes that the prevailing reference point is the target return, even when the risk-free rate becomes negative. To check this conjecture, we need to compare the variations in the proportion of money invested in the risky asset when the risk-free rate drops from 0% to -2% in the Low-Target-Return condition (from point B to C on Figure 1), with the variations in the proportion of money invested in the risky asset when the risk-free rate drops from 4% to 2% in the High-Target-Return condition (from point G to H on Figure 1). Since the risk-free rate is below the target return and the distance between the risk-free rate and the target return is the same in both conditions, participants should exhibit similar risk-seeking behavior if the target return still prevails as reference point when the risk-free rate drops below zero. This means that the amount invested in the risky asset should increase in the same proportion in both conditions. Table 8 reports the results of a two-sample t-test addressing the difference between the average variations in risk-taking in the High- and Low-Target-Return conditions. This difference of means is positive (1.65%) but insignificant at any conventional level. This indicates that the average marginal increase in risk-taking is similar in both conditions, which supports our last hypothesis. The target return still acts as the prevailing reference point for our participants who face a negative risk-free rate. Having the risk-free rate crossing the zero-level does not especially affect their increase in risk-taking ; they continue to adjust their allocation to the risky asset based on the distance to their target return. We conclude that the move of the risk-free rate into negative territory does not activate any additional (i.e., extra) risk-taking.

Table 8: Variations in risk-taking and NIR

Δ when R_F moves from 0% to -2%	6.00%***
Δ when R_F moves from 4% to 2%	4.35%***
Difference	1.65%
t-value	0.20
p-value	0.8387

This table reports the average variations in the allocation to the risky asset when the risk-free rate moves from 0% to -2% in the Low-Target-Return condition and from 4% to 2% in the High-Target-condition, the difference between these average variations, and the corresponding t-value and p-value from a two-sample t-test. *** indicates when the average variations statistically differ from zero (at the 1% level).

4.3 Multivariate analysis

To further investigate the impact of the target return on the risk-taking behavior, we estimate an OLS regression model wherein the dependent variable, $Risk_taking_{i,d}$, is the proportion of money invested in the stock index fund by participant i at decision d . The set of explanatory variables includes both individual-varying variables as well as decision-varying variables.

The individual-varying variables include some sociodemographics that are usual control variables (age, gender, student status) and the additional individual characteristics collected at the end of the experiment.

For gender, we use a binary variable ($Male_i$) set to one for males. We also use a binary variable for student status ($Student_i$) that is set to one for participants who are students (and not employed). Risk tolerance ($Risk_tol_i$) is the participant's average score based on the answers to the six questions taken from the DOSPERT scale, which ranges from 0 (extremely unlikely to take a risk) to 7 (extremely likely to take a risk). We use binary variables to characterize participants who own a real deposit account (Dep_Acc_i is set to one) and an actual portfolio of financial assets ($Stocks_i$ is set to one). To control for the condition to which the participant has been assigned, we add another binary variable (Low_cond_i) set to one for participants in the Low-Target-Return condition.

The decision-varying variables aim at capturing the possible levels for the risk-free rate. We consider the base case when the risk-free rate matches the target return, i.e., 6% (2%) in High- (Low-) Target-Return condition, respectively. Hence, $Below_target_{i,d}$ is a binary variable set to one when the risk-free rate falls short of target return. Similarly, $Above_target_{i,d}$ is another binary variable set to one when the risk-free rate is above the target return. This leads to the following baseline model:

$$Risk_taking_{i,d} = \beta_0 + \beta_1 Low_cond_i + \beta_2 Age_i + \beta_3 Male_i + \beta_4 Student_i + \beta_5 Risk_tol_i + \beta_6 Dep_Acc_i + \beta_7 Stocks_i + \beta_8 Below_target_{i,d} + \beta_9 Above_target_{i,d} + \epsilon_{i,d} \quad (1)$$

We also estimate an alternative to the baseline model, wherein we split the two decision-varying variables into the four following dummies: $Below_4\%_{i,d}$, $Below_2\%_{i,d}$, $Above_2\%_{i,d}$, $Above_4\%_{i,d}$, which are set to one when the risk-free rate is 4% below, 2% below, 2% above, and 4% above the target return, respectively. The resulting model is given as follows:

$$Risk_taking_{i,d} = \beta_0 + \beta_1 Low_cond_i + \beta_2 Age_i + \beta_3 Male_i + \beta_4 Student_i + \beta_5 Risk_tol_i + \beta_6 Dep_Acc_i + \beta_7 Stocks_i + \beta_8 Below_4\%_{i,d} + \beta_9 Below_2\%_{i,d} + \beta_{10} Above_2\%_{i,d} + \beta_{11} Above_4\%_{i,d} + \epsilon_{i,d} \quad (2)$$

We provide the results in Table 9, wherein Model 1 corresponds to our baseline model (Eq. 1), Model 2 to the alternative (Eq. 2), and Model 3 is just Model 2 with an interaction variable ($Low_cond_i * Below_4\%_{i,d}$) that refers to situations when participants in the Low-Target-Return condition face NIR.

Looking at the decision-varying variables in our baseline model (Model 1), the expected risk-seeking behaviors emerge. When the risk-free rate is below the target return, participants' marginal increase in risk-taking is 22% higher, compared to situations when the risk-free rate just matches the target return. However, when the risk-free rate is above the target return, the coefficient estimate is positive but not significant, which does not provide evidence of real change in risk-taking. Next, the coefficient estimate of the variable controlling for the condition to which participant have been assigned is positive but not significant. This means that participants exhibit similar risk-taking behaviors in both conditions. In other words, either High- or Low-Target-Return condition does not induce higher or lower risk-taking per se.

Focusing on the individual-varying variables in our baseline model (Model 1), the relationship between risk-taking and risk tolerance is significantly negative, which is rather counter-intuitive. All else being equal, when risk tolerance increases by one unit, the allocation to the stock index fund decreases by about 4%. We could relate this findings to the literature showing that some inconsistencies might exist

between subjective risk tolerance (i.e., risk tolerance assessed through a questionnaire) and objective risk tolerance (i.e., risk tolerance assessed through actual behavior). Although Chang et al. (2004) find a positive correlation between subjective and objective risk tolerance, Jianakoplos (2011) reports contrasting results. Using survey data to confront statements with actual behaviors, this author finds that individuals' one-third of total wealth is invested in risky assets, although they stated an unwillingness to take risks. Mazzoli et al. (2017) find that such inconsistencies are more pronounced for investors with low financial literacy, high income, and no children. As for age, the relationship with risk-taking is significantly negative. When participants get one year older, their allocation to the stock index fund decreases by almost 0.4%. As far as gender is concerned, we find that males invest around 9% more in the stock index fund, compared to females. This finding is consistent with the literature providing evidence that males are more risk-takers than females (e.g., Sung and Hanna (1996), Jianakoplos (2011)). Finally, students appear to be overall more risk-averse than employees, i.e., the allocation to the stock index fund is about 10% lower when participants are students. The owning of a real deposit account or an actual portfolio of financial assets does not help explain heterogeneity in risk-taking across participants.

The results of our alternative model (Model 2) are totally consistent with those of the baseline model (Model 1). Focusing on the four dummies related to the position of the risk-free rate vis-à-vis the target return, risk-seeking behaviors are confirmed when the risk-free rate is below the target return. Specifically, when the risk-free rate is 4% below the target return, participants invest almost 23% more in the stock index fund, compared to when the risk-free rate just matches it. Similarly, when the risk-free rate is 2% below the target return, the allocation to the risky asset decreases by about 20%. These results are both statistically and economically significant. All the other variables in Model 2 display coefficient estimates that are very close and qualitatively similar to those exhibited in Model 1.

The main interest of the third model (Model 3) in Table 9 is to include an interaction variable ($Low_cond_i * Below_4\%_{i,d}$) that refers to situations when participants in the Low-Target-Return condition face NIR. The coefficient estimate of this interaction term is positive but not statistically significant. This confirms that the zero threshold does not overtake the role of reference point when the risk-free rate becomes negative. It is the target return that still plays this crucial role and drives risk-taking.

In the aforementioned panel data regression models, the dependent variable is defined as the proportion of money allocated to the risky asset. Technically, this means that our dependent variable is bounded between 0 and 1, which may question the use of an OLS regression. As a robustness check, we replicate our approach using three fractional logit models, which are econometrically more appropriate for dependant variables taking values raging from zero to one. The results, which are qualitatively similar for our main variables of interest, are available in in Appendix D.

Table 9: Panel data regression models for risk-taking behavior

Parameter	Model 1	Model 2	Model 3
Intercept	62.2774***	62.5732***	63.5022***
Low_cond	4.6434	4.7395	2.7544
Age	-0.3871**	-0.3858**	-0.3861**
Male	8.8630***	8.9578***	8.9388***
Student	-10.1277***	-9.9914***	-10.0187***
Risk_tol	-4.4225***	-4.3316***	-4.3498***
Dep_Acc	2.2123	2.1062	2.1275
Stocks	5.3706	5.1874	5.2241
Below_target	22.4822***		
Above_target	3.6799		
Below_4%		22.6170***	18.0914***
Below_2%		20.4369***	20.5357***
Above_2%		4.5191	4.5620
Above_4%		1.7532	1.7961
Low_cond*Below_4%			9.8294
N	405	405	405
R ²	0.1501	0.1477	0.1515

This table reports the results for panel data regression models wherein the dependent variable, $Risk_taking_{i,d}$, is the proportion of money invested in the stock index fund by participant i at decision d . The individual-varying variables include sociodemographics that are usual control variables and additional individual characteristics. $Male_i$ is a binary variable set to one for males. $Student_i$ is a binary variable set to one for participants who are students (and not employed). $Risk_tol_i$ refers to the participant's average score based on the answers to the six questions taken from the DOSPERT scale, which ranges from 0 (extremely unlikely to take a risk) to 7 (extremely likely to take a risk). Dep_Acc_i and $Stocks_i$ are binary variables set to one for participants who own a real deposit account and an actual portfolio of financial assets, respectively. Low_cond_i is another binary variable set to one for participants in the Low-Target-Return condition. Among the decision-varying variables in Model 1, $Below_target_{i,d}$ is a binary variable set to one when the risk-free rate falls short of target return, while $Above_target_{i,d}$ is another binary variable set to one when the risk-free rate is above the target return. As alternatives in Model 2, the four following dummies are used: $Below_4\%_{i,d}$, $Below_2\%_{i,d}$, $Above_2\%_{i,d}$, $Above_4\%_{i,d}$, which are set to one when the risk-free rate is 4% below, 2% below, 2% above, and 4% above the target return, respectively. In Model 3, $Low_cond_i * Below_4\%_{i,d}$ is an interaction term that refers to situations when participants in the Low-Target-Return condition face NIR. ***, **, and * indicate significance at 1%, 5%, and 10%, respectively. N gives the number of observations (i.e., decisions).

5 Conclusion

This paper contributes to the growing body of literature on what stimulates the appetite of individuals for risk when facing NIRs. The novelty of our approach is to introduce a variable of particular interest in investment decisions, that is the target return. Unlike past research that considers the zero threshold as the reference point by default (Bracha, 2020; Ganzach and Wohl, 2018; Lian et al., 2018; Baars

et al., 2020; David-Pur et al., 2020), we run an experiment wherein explicit target returns are randomly attributed to participants. Building on both Heath et al. (1999) and Larrick et al. (2009) who postulate that goals incorporate the three main principles of Prospect Theory (i.e., reference point, loss aversion, and diminishing sensitivity), we posit that target returns act as goals that alter the psychological value of outcomes and impact the appetite for risk. Specifically, we first check whether the target return serves as a reference point and effectively impacts the risk-taking behavior of participants when making investment decisions. Second, we determine which of the target return or the zero threshold is the prevailing reference point when the risk-free rate drops below zero. Our bespoke experiment is a mixed design with one between - the target return condition (High vs. Low) and two within factors - the interest rate on the deposit account (i.e., risk-free rate) and the stock index fund return (i.e., the risky asset return with a risk premium kept constant at 5%).

Our original approach provides insightful results about the "reach for yield" among individuals. Overall, our findings confirm that the target return is the reference point of participants when making investment decisions. This holds whatever the level of the target return, i.e., in both High- and Low-Target-Return conditions. First, participants invest more in the stock index fund when the risk-free interest rate is below their target return than when it is above their target return. Second, consistent with loss aversion, participants' reactions when the risk-free rate falls below their target return (i.e., increase in risk-taking) exceed the corresponding reactions when the risk-free rate goes above their target return. Third, we provide partial evidence of diminishing sensitivity, which predicts a greater impact of a given shift in the risk-free rate when occurring near the target return (than far from it). Participants exhibit the expected behavior, but only when the risk-free rate fluctuates below their target return.

Our result for the specific case when the risk-free rate drops below zero in the Low-Target-Return condition delivers new insights. Comparing the average increases in risk-taking for similar shifts of the risk-free rate in both conditions (i.e., from 0% to -2% in the Low-Target-Return condition and from 4% to 2% in the High-Target-Return condition), we find no significant difference. This indicates that the target return still acts as the prevailing reference point for our participants who face a negative risk-free rate. Having the risk-free rate crossing the zero-level does not especially affect risk-taking when people have an explicit target return (that differs from zero); people continue to adjust their allocation to the risky asset based on the distance to their target return. The larger the distance between the target return and the risk-free rate, the lower is the marginal increase in risk-taking. Hence, the move of the risk-free rate into negative territory does not activate any additional (i.e., extra) risk-taking.

Our additional findings based on panel data regression models confirm the expected risk-seeking behaviors when the risk-free rate is below the target return. Whatever the target return, when the risk-free rate is 2% (4%) below the target return, participants decrease their allocation to the risky asset by about 20% (23%). Such reductions are both statistically and economically significant. Furthermore, when focusing on situations when participants in the Low-Target-Return condition face a negative risk-free rate, our multivariate results confirm that the target return is still the main driver of risk-taking as prevailing reference point.

Our results show that the zero-level does not act as a reference point when a target return is ex-

plicitly set for investment decisions, which provides useful insights for both monetary policymakers and practitioners in the financial industry. Indeed, it reveals that NIRs is a limited tool to stimulate the appetite for risk since the risk-taking behavior of individuals is mostly driven by their target return when making investment decisions. Unlike in our experiment wherein each participant has been assigned with a given target return, regular people are likely to have in mind some return objectives when making real investment decisions. More importantly, return objectives (and more generally financial objectives) are key information that investment firms are used to rely on for portfolio management. Such an item is usually covered in Investment Policy Statements (IPS),¹⁷ meaning that investors are led to make their target return explicit. Hence, if their target return differs from zero (which is highly plausible), these investors are expected to increase their risk-taking under NIRs, but only in a marginal way. The higher the target return, the smaller will be the marginal increase in risk-taking when the risk-free rate turns out negative (since the shift is going to happen far from their reference point). Put differently, investors who are expected to increase the most their risk-taking are those who have low target returns (i.e., close to zero). The limitation of NIRs as an extra booster of risk-taking is particularly relevant in the European Union member states wherein the MiFID regulation¹⁸ has made mandatory the Suitability questionnaire,¹⁹ which is a kind of regulated IPS wherein each retail client is asked to declare his/her return objective (i.e., make his/her target return explicit) before investing.

Our experimental findings, as ecologically valid as they may be, do not fully reflect the dynamics and complexity of the investment decisions people make in reality. One shortcoming in our approach is related to how we manipulate the between-subject factor, i.e., we assign given target returns to participants. As mentioned above, in reality, people are not assigned with an explicit target return. They are rather asked to declare it, either to their financial intermediary or in a regulated IPS (such the Suitability MiFID questionnaire). To further investigate the role of the target return on risk-taking, it would be relevant to examine how individuals make investment decisions while setting their own target returns. Another limitation in our experiment is that leverage is not allowed, meaning that participants could not borrow money at the risk-free rate to invest more than their initial endowment in the risky asset. In the future, follow-up studies might be designed to overcome such shortcomings.

¹⁷The following items are primordial for an IPS used in portfolio management delegation: goals and objectives, fears and concerns, investment time frame, expected outside mortality, retirement time frame, shorter-term financial needs, risk tolerance attitudes, and risk capacity (Boone and Lubitz, 2004)

¹⁸MiFID stands for Markets in Financial Instruments Directive. This piece of regulation came into force in November 2007 in European Union (EU) member states. MiFID I (2004/39/EC) is the first version of this piece of regulation that has been updated as MiFID II in January 2018 (known as MiFID II (2014/65/UE)). The European Commission website provides a complete description of this MiFID regulation. See https://ec.europa.eu/info/law/markets-financial-instruments-mifid-ii-directive-2014-65-eu_en.

¹⁹MiFID distinguishes three types of services: execution of orders, financial advice, and portfolio management. Investors who require financial advice or portfolio management have to complete the *Suitability* test, which aims at ensuring that the instruments and services offered by the investment firm meet the investor's knowledge and experience, his/her investment objectives as well as his/her financial capacity.

References

- Agarwal, R. and M. Kimball (2015). Breaking Through the Zero Lower Bound. Technical report, IMF Working Paper WP/15/224. Washington: International Monetary Fund, October.
- Altavilla, C., L. Burlon, M. Giannetti, and S. Holton (2019). Is there a zero lower bound? The effects of negative policy rates on banks and firms. ECB Working Paper, No. 2289.
- Baars, M., H. Cordes, and H. Mohrschladt (2020). How negative interest rates affect the risk-taking of individual investors: Experimental evidence. *Finance Research Letters* 32, 101179.
- Blais, A.-R. and E. Weber (2006, 02). A Domain-Specific Risk-Taking (DOSPERT) scale for adult populations. *Judgment and Decision Making* 1, 33–47.
- Boone, N. M. and L. S. Lubitz (2004). *Creating an Investment Policy Statement: Guidelines & Templates*. AdvisorPress.
- Boungou, W. (2020). Negative interest rates policy and banks’ risk-taking: Empirical evidence. *Economics Letters* 186, 108760.
- Bracha, A. (2020). Investment Decisions and Negative Interest Rates. *Management Science* 66(11), 5316–5340.
- Brown, B. (2018). Negative Interest Rates and the War Against Cash. In *The Case Against 2 Per Cent Inflation*, pp. 149–163. Springer.
- Chang, C.-C., S. A. DeVaney, and S. T. Chiremba (2004). Determinants of subjective and objective risk tolerance. *Journal of Personal Finance* 3(3), 53–67.
- Chen, D., M. Schonger, and C. Wickens (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88–97.
- David-Pur, L., K. Galil, and M. Rosenboim (2020, 06). To Decrease or Not to Decrease: The Impact of Zero and Negative Interest Rates on Investment Decisions. *Journal of Behavioral and Experimental Economics* 87, 101571.
- Ganzach, Y. and A. Wohl (2018, 07). A Behavioral Theory of the Effect of the Risk-Free Rate on the Demand for Risky Assets. *Journal of Behavioral and Experimental Economics* 76, 23–37.
- Heath, C., R. P. Larrick, and G. Wu (1999). Goals as Reference Points. *Cognitive Psychology* 38(1), 79–109.
- Jianakoplos, N. (2011, 02). Invest As I Say, Not As I Do? Gender Differences In Financial Risk Preferences. *International Business Economics Research Journal (IBER)* 1.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An Analysis of Decision under Risk. *Econometrica* 47(2), 263–291.

- Kimball, M. S. (2015). Negative Interest Rate Policy as Conventional Monetary Policy. *National Institute Economic Review* 234(1), R5–R14.
- Larrick, R. P., C. Heath, and G. Wu (2009). Goal-Induced Risk Taking in Negotiation and Decision Making. *Social Cognition* 27(3), 342–364.
- Leung, W. S. and Z. Zhu (2018). Zero Lower Bound and Investors’ Risk Taking on Fixed Income Mutual Funds. Available at SSRN: <https://ssrn.com/abstract=3103496> or <http://dx.doi.org/10.2139/ssrn.3103496>.
- Lian, C., Y. Ma, and C. Wang (2018). Low Interest Rates and Risk-Taking: Evidence from Individual Investment Decisions. *The Review of Financial Studies* 32(6), 2107–2148.
- Locke, E. A. and G. P. Latham (1990). *A theory of goal setting & task performance*. Prentice-Hall, Inc.
- Maggio, M. D. and M. Kacperczyk (2017). The unintended consequences of the zero lower bound policy. *Journal of Financial Economics* 123(1), 59 – 80.
- Mazzoli, C., N. Marinelli, and F. Palmucci (2017, 05). Mind the Gap: Inconsistencies Between Subjective and Objective Financial Risk Tolerance. *Journal of Behavioral Finance* 18, 1–12.
- Payne, J. W., D. J. Laughhunn, and R. Crum (1980). Translation of Gambles and Aspiration Level Effects in Risky Choice Behavior. *Management Science* 26(10), 1039–1060.
- Rechenberg, T., D. Gutt, and D. Kundisch (2016). Goals as Reference Points: Empirical Evidence from a Virtual Reward System. *Decision Analysis* 13, 127–151.
- Shefrin, H. and M. Statman (2000, 06). Behavioral Portfolio Theory. *Journal of Financial and Quantitative Analysis* 35, 127–151.
- Sung, J. and S. D. Hanna (1996). Factors Related to Risk Tolerance. *Capital Markets eJournal*.
- Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297–323.

Appendices

A Probability weighting function




Figure 4: Probability weighting function

This figure presents the probability weighting function for gains and losses. The horizontal axis refers to the probabilities, while the vertical axis refers to corresponding decision weights. The grey (orange) line corresponds to the decision weights related to probabilities of losses (gains), while the blue line follows the expected utility theory wherein decision weights match probabilities.

B Domain-Specific-Risk-Taking (DOSPERT) scale - Statements

For each of the following statements, please indicate the likelihood that you would engage in the described activity or behavior if you were to find yourself in that situation. Provide a rating from Extremely Unlikely to Extremely Likely.

Statement 1: Betting a daily income at the horse races.

Statement 2: Investing 10% of your annual income in a moderate growth diversified fund.

Statement 3: Betting a daily income at a high-stake poker game.

Statement 4: Investing 5% of your annual income in a very speculative stock.

Statement 5: Betting a daily income on the outcome of a sporting event.

Statement 6: Investing 10% of your annual income in a new business venture.

C Analysis of target returns in the real world

Table 10: Difference in target returns in reality between conditions

High-Target-Return	19.53%
Low-Target-Return	10.63%
Difference (High - Low)	8.90%
t-value	3.93
p-value	<0.0001

This table presents the average target returns participants would set in reality in both conditions, the difference between these target returns, the corresponding t-values, and p-values from a two-sample t-test.

D Results based on Fractional logit models

Table 11: Fractional logit regressions for risk-taking behavior

Parameter	Model 1		Model 2		Model 3	
	Estimate	OR	Estimate	OR	Estimate	OR
Intercept	0.5378	1.7122	0.5481	1.7300	0.5910	1.8058
Low_cond.	0.2035	1.2257	0.2059	1.2286	0.1185	1.1258
Risk_tol	-0.1923*	0.8251	-0.1892*	0.8276	-0.1906*	0.8265
Dep_Acc	0.0967	1.1015	0.0935	1.0980	0.09531	1.1000
Stock	0.2338	1.2634	0.2280	1.2561	0.2290	1.2573
Age	-0.01670	0.9834	-0.01663	0.9835	-0.01665	0.9835
Male	0.3853*	1.4701	0.3873*	1.4730	0.3880*	1.4740
Student	-0.4359*	0.6467	-0.4310*	0.6499	-0.4324*	0.6489
Below_target	0.9570***	2.6039				
Above_target	0.1514	1.1635				
Below_4%			0.9741***	2.6488	0.7621***	2.1428
Below_2%			0.8722***	2.3922	0.8737***	2.3958
Above_2%			0.1892	1.2083	0.1900	1.2092
Above_4%			0.0757	1.0787	0.07666	1.0797
Low_cond.*Below_4%					0.4884	1.6297
N	405		405		405	

This table reports the results for three fractional logit regression models wherein the dependent variable, $Risk_taking_{i,d}$, is the proportion of money invested in the stock index fund by participant i at decision d . The individual-varying variables include sociodemographics that are usual control variables and additional individual characteristics. $Male_i$ is a binary variable set to one for males. $Student_i$ is a binary variable set to one for participants who are students (and not employed). $Risk_tol_i$ refers to the participant's average score based on the answers to the six questions taken from the DOSPERT scale, which ranges from 0 (extremely unlikely to take a risk) to 7 (extremely likely to take a risk). Dep_Acc_i and $Stocks_i$ are binary variables set to one for participants who own a real deposit account and an actual portfolio of financial assets, respectively. Low_cond_i is another binary variable set to one for participants in the Low-Target-Return condition. Among the decision-varying variables in Model 1, $Below_target_{i,d}$ is a binary variable set to one when the risk-free rate falls short of target return, while $Above_target_{i,d}$ is another binary variable set to one when the risk-free rate is above the target return. As alternatives in Model 2, the four following dummies are used: $Below_4\%_{i,d}$, $Below_2\%_{i,d}$, $Above_2\%_{i,d}$, $Above_4\%_{i,d}$, which are set to one when the risk-free rate is 4% below, 2% below, 2% above, and 4% above the target return, respectively. In Model 3, $Low_cond.*Below_4\%_{i,d}$ is an interaction term that refers to situations when participants in the Low-Target-Return condition face NIR. ***, **, and * indicate significance at 1%, 5%, and 10%, respectively. 'OR' refers to odds ratios. N gives the number of observations (i.e., decisions).