Matteo Benuzzi, Matteo Ploner

Skewness-seeking behavior and financial investments

CEEL Working Paper 1-23

Cognitive and Experimental Economics Laboratory

Via Inama, 5 38100 Trento, Italy

http://www-ceel.economia.unitn.it
tel. +39.461.282313

Skewness-seeking behavior and financial investments

Matteo Benuzzi[a*], Matteo Ploner[b]

May 30, 2023

Abstract

Recent theoretical and empirical contributions have demonstrated the significance of higher-order moments, such as skewness, in influencing financial decisions. Most current experimental literature relies on lotteries with a limited number of potential outcomes, which do not accurately represent real-life investments. To address this gap, we conducted a pre-registered experiment that examines preferences toward investment opportunities with varying skewness using continuous distributions. Our findings reveal several key insights. Firstly, there is an overall preference for positively skewed distributions of outcomes. Secondly, we observed a substitution effect between risk-taking, as measured by variance, and the direction of skewness. Lastly, we established a positive correlation between skewness-seeking behavior and speculative behavior and a negative correlation between skewness-seeking behavior and risk perception of positive skewness.

Keywords Skewness · Risk-taking · Stochastic Dominance · Experiment

JEL Classification C91 · D81 · G11

---

[*] Corresponding author. Email: matteo.benuzzi@unitn.it

[a] School of Social Sciences, University of Trento
[b] Department of Economics and Management, University of Trento

1 Introduction

Skewness is a measure of the asymmetry of a distribution, such as the distribution of returns on investment. While the literature has traditionally focused on the first two moments of the distribution - expected return and variance - higher-order moments have increasingly been studied from multiple perspectives. Skewness, which is usually measured with the third standardized moment, has been indeed associated with several phenomena, effects, and anomalies, such as the long-shot anomaly on the horse track (Golec and Tamarkin, 1998) and in online lotteries (Garrett and Sobel, 1999), the volatility smile (Barberis and Huang, 2008; Boyer and Vorkink, 2014), the preference for lottery-like stocks (Boyer et al., 2010), the underperformance of IPOs (Barberis, 2013), the underperformance of high-skewness stocks (Amaya et al., 2015), and the conglomerate discount (Schneider and Spalt, 2016). The common denominator of all these phenomena is that some individuals who find the combination of low probabilities and large outcomes particularly attractive overpay for access to these investments/gambles, which, as a result, tend to yield lower returns. In this context, skewness preference is a concept used to refer to the preference for a positively skewed distribution of outcomes - i.e., one with a long right tail - over a negatively skewed one.

The experimental literature includes several contributions concerning the topic of skewness preferences. Most of the literature has examined preferences over skewed distributions of outcomes using binary and three-outcome lotteries, departing significantly from continuous distributions, which are more appropriate representations of the outcomes that investors face on the markets. The results of these experiments may not be easily generalized to continuous distributions due to the presence of biases and heuristics that may affect judgment (Holzmeister et al., 2020; Summers and Duxbury, 2006; Vrecko et al., 2009).

We designed and conducted an experiment to investigate the role of skewness in financial investments using continuous distributions of outcomes. Consistent with some of the existing literature, we found evidence of preferences for outcome distributions with a positive skewness coefficient. Furthermore, we identified two channels of skewness preferences: the first is related to speculative behavior, and the second is related to risk perception. On the contrary, risk preferences do not seem to play a role. Finally, contrary to most literature, we found that a negatively skewed environment encourages more risk-taking than a positively skewed one: subjects forced (by treatment) to choose a negatively skewed distribution chose distributions with a higher standard deviation than those forced to choose a positively skewed one.

A better understanding of the role of skewness in ordering alternative investment perspectives is essential for advancing our understanding of decision-making under risk and may also have important implications for investment decisions in the real world. For example, considering the third distribution moment, in addition to the traditional first two, may open new opportunities to engineer financial products that better suit investors' needs and preferences. New products of this kind may represent an opportunity for subscribers of financial products and allow for a more efficient risk distribution in the financial system.

The remainder of the paper is organized as follows. In section two, we summarize some of the existing literature on the role of skewness in financial decisions, with a focus on the experimental literature; in section three, we describe the experimental design and the research questions; in section four, we report the main results; in section five, we discuss the main findings; finally, in section six we report our conclusions.

2 Literature

2.1 The role of skewness from a theoretical perspective

Skewness is a measure of the asymmetry of a distribution, and it is usually measured with the coefficient of skewness, which is equal to the third standardized moment: $\tilde{\mu}_3 = E\left[\left(\frac{X-\mu}{\sigma}\right)^3\right]$. Preference for skewness is a concept used to refer to the preference for a positively skewed distribution over a negatively skewed one. The traditional approach relating the third moment with expected utility theory suggests that a higher skewness coefficient should increase utility (Arditti, 1967): assuming positive marginal utility of wealth, risk aversion, and decreasing local risk aversion (Pratt, 1964), implies that $U^{'''}(w) > 0$. This would mean that the higher the skewness coefficient, the higher the utility. However, this approach neglects all the other higher moments of the distribution of outcomes, and for this reason, it has been criticized by Brockett and Kahane (1992) and Brockett and Garven (1998), who proved that, under EUT, utility is not necessarily increasing in skewness.[1]

An individual with a positive third-order derivative of the utility function, $U^{'''}(w) > 0$, is said to be prudent, a concept introduced by Kimball (1990). Eeckhoudt and Schlesinger (2006) provided a model-free definition of prudence which is equivalent to $U^{'''}(w) > 0$ in the EUT framework: given the ES pair of lotteries $B_3 = [-k, \epsilon]$ and $A_3 = [-k + \epsilon, 0]$, preference for the former over the latter implies prudence[2]. By preferring B_3, a prudent decision-maker prefers to add the zero-mean risk ϵ to the positive state where she earns zero over the negative state where she earns the sure loss k. Prudence has indeed been found to be prevalent in several experimental works (Deck and Schlesinger, 2010; Ebert and Wiesen, 2011; Fairley and Sanfey, 2020; Heinrich and Shachat, 2020; Noussair et al., 2014).

Skewness-seeking behavior is defined as the preference for a distribution of outcomes with a higher skewness coefficient over another with the same expected value, variance, and kurtosis (Ebert and Wiesen, 2011). From a theoretical perspective, a prudent individual should be a skewness-seeker, but the opposite may not necessarily be true. Ebert and Wiesen (2011) provided experimental evidence in support of this.

Risk is traditionally measured with volatility (i.e., standard deviation). Still, a shortcoming of volatility is the equal treatment that all deviations from the mean receive, regardless of whether they are positive or negative (Markowitz, 1991). Holzmeister et al. (2020) claim that the probability of experiencing a loss is the driver of risk perception of skewness, and therefore positively skewed distributions would be perceived riskier. More in general, the use of probabilities to experience losses below some threshold may lead to different rankings of distributions differing in skewness depending on the chosen threshold as well as the characteristics of the distributions. Finally, risk measures such as the behavioral risk measure σ_B^2 introduced by Davies and De Servigny (2012) directly incorporate higher-order mo-

[1] The reason is that when the skewness of a lottery changes, the other moments may also change. The truncation of the Taylor series at the third term neglects all these other moments, ignoring a portion of the true utility associated with the lottery.

[2] For both lotteries each state has a 50% probability, k is a positive constant and ϵ is a zero-mean random variable

ments. In particular, σ_B^2 is decreasing in skewness (keeping variance and kurtosis constant).

2.2 The role of skewness in the experimental literature

One of the first studies of skewness preferences was Mao's (1970): he asked managers to decide between two binary lotteries with the same mean and variance but differing skewness. Managers were almost equally split between the alternatives if the investment represented a small portion of the company resources. However, when the lotteries represented the whole outcome of the company, they all picked the positively skewed distribution because of its better downside (i.e., a better outcome in the negative state). Brünner et al. (2011) studied skewness-seeking behavior using pairs of binary prospects with the same mean and variance but differing skewness and once again found evidence of skewness-seeking behavior, with about 60% of the subjects choosing the prospect with larger skewness more than half of times. Ebert and Wiesen (2011) used Mao lottery pairs (i.e., sets of two binary lotteries with the same mean, variance, and kurtosis, but different skewness) and found skewness-seeking behavior for about 75% of the subjects. Åstebro et al. (2015) tested skewness preferences with a variation of Holt and Laury (2002)'s multiple price list format, and found that subjects made, on average, skewness-seeking choices. Similarly, Ebert (2015) found that 64% of subjects make skewness-seeking choices, and Dertwinkel-Kalt and Köster (2020) found a preference for positive skewness.

Grossman and Eckel (2015) used a variation of Eckel and Grossman lotteries (2002; 2008), and found that more than 80% of the subjects were skewness-seekers. Taylor (2020) used a slight modification of Grossman and Eckel's (2015) design aimed at reducing the impact that loss aversion may have had on skewness-seeking behavior and indeed found that while this behavior was still prevalent, it was less frequently observed compared to Grossman and Eckel's study. Unlike the previous studies, both Grossman and Eckel (2015) and Taylor (2020) did not use binary prospects but three outcome prospects. Bougherara et al. (2021) found mostly skewness-avoidance in an experiment where they elicited certain equivalents of three-outcome prospects[3]. In a subsequent study with similar experimental settings, Bougherara et al. (2022) found mostly skewness-seeking behavior.

While using binary prospects seems to lead to skewness-seeking behavior, introducing other outcomes makes the situation less clear-cut. The impact of a shift from finite outcome prospects to continuous distributions is even more dramatic. Vrecko et al. (2009) found that the form used to represent an investment affects the decision: subjects were found to be skewness-seekers when a cumulative distribution function was used to represent the alternatives, whereas they were skewness-avoiders when a probability density function was utilized instead. Skewness-avoidance in the *density treatment* was rooted in anchoring (Tversky and Kahneman, 1974): the peak of the distribution would serve as an anchor for the estimation of the unknown mean (Summers and Duxbury, 2006). As a result, when the expected value of the distribution is unknown, it is overestimated for negatively skewed distributions and underestimated for positively skewed ones (Vrecko et al., 2009). Figure 1 shows that as the skewness coefficient increases, the mode of the distribution decreases (moving to the left), even if the expected return is the same for all three distributions. Thus, using probability density functions may discourage skewness-seeking behavior due to biased risk perception and incorrect estimation of the expected return. Holzmeister et al. (2020) elicited risk perception and investment propensity of continuous distributions differing in skewness, represented by histograms of samples from such distributions. They found that positively skewed distributions are generally perceived as riskier than negatively skewed ones by financial professionals and laypeople, with this phenomenon being driven by the higher probability of a loss. Likewise, investment propensity was negatively associated with risk perception, and thus, positively skewed distributions of outcomes were less likely to be chosen.

3 Methods

While the literature suggests that positive skewness should, at least to some extent, be associated with a lower level of risk, continuous distributions seem to lead sub-

[3]They found that subjects prefer highly negative skewed prospects over low negative skewed ones, and low positively skewed prospects over highly positive skewed ones, both for high and low variance. However, they preferred low positively skewed prospects over low negatively skewed prospects. Prospects had the same mean, variance, and kurtosis but differed in skewness.

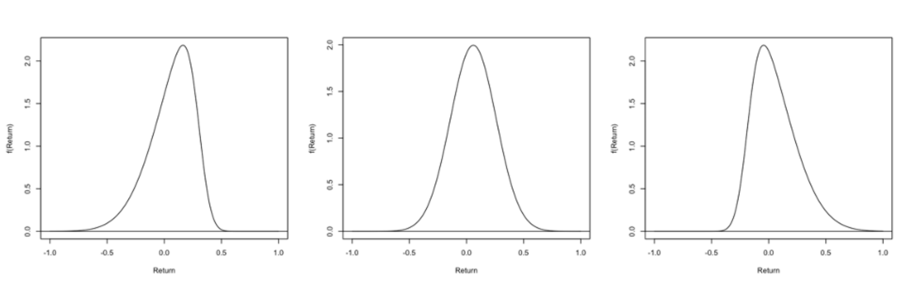

Figure 1: Probability density function of three distributions. From the left to right, the first is negatively skewed, the second is symmetric, and the third is positively skewed. The distributions have the same expected return and variance.

jects to believe the opposite. Moreover, among the studies reported in the previous section, some focus on decisions in a gains-only framework (Åstebro et al., 2015; Bougherara et al., 2022; Brünner et al., 2011), and some consider only positively skewed and symmetric distributions of outcomes (Åstebro et al., 2015; Grossman and Eckel, 2015; Taylor, 2020). Finally, all studies are characterized by significant heterogeneity in skewness preferences. For these reasons, the topic deserves further attention in a comprehensive framework: our pre-registered experimental design[4] aims at studying the preferences over skewed continuous distributions of outcomes, considering both gains and losses, as well as positive, zero and negative levels of skewness.

3.1 Research questions

The experiment aims to study skewness preferences under several perspectives using continuous distributions of outcomes. Within our experimental setting, we address four main research questions.

The first research question concerns the relationship between skewness and risk-taking: *"How do skewness and risk-taking interact?"*. We address this question in rounds seven and eight (see section 4.1).

The second research question concerns skewness preferences: *"Do subjects exhibit skewness-seeking behavior when opportunities are shown using continuous distributions of outcomes?"*. We address this question in rounds one, four, five, and

[4] Link: `https://osf.io/9q72b/?view_only=9b5327eedc4b46d187e34568d52d9f48`.

six (see sections 4.2 and 4.3).

The third research question regards the trade-off between skewness and expected return: *"Do subjects trade off skewness with expected returns, or do they exhibit mean-variance preferences?"*. We address this question in rounds two and three (see section 4.4).

The fourth research question is about the drivers of skewness preferences: *"What are the characteristics driving skewness preferences?"*. We address this question by combining choices of rounds one to six with measures collected outside part one of the experiments, which concern speculative behavior, risk, and loss preferences (see section 4.5).

3.2 Experimental design

The online experiment was programmed and executed with oTree (Chen et al., 2016). We present here the different parts of the experiment (see the appendix for a detailed description). After an introductory non-incentivized part, participants were given a tutorial on the experimental framework. The tutorial reviewed the concepts of probability density function, variance, and skewness and offered the chance to see how changes in these moments would affect the distribution of outcomes. The tutorial was followed by a mandatory comprehension check, and then subjects were asked about their general aspirations about investments. Then the first part of the experiment began: subjects played eight rounds, making incentivized decisions over distributions differing in skewness. In the second part of the experiment, subjects' risk preferences were elicited using a modified version of Holt and Laury (2002) multiple price list (the payoffs of the low-risk lottery A were £0.6 and £0.8, while the payoffs of high-risk Lottery B were £0.1 and £1.2), with one of the ten choices randomly selected for payment. Finally, a demographic and a financial-behavior questionnaire were administered.

Like Ebert (2015), we framed our alternative distributions in terms of returns rather than outcomes: the final payment from part one of the experiment was computed compounding the returns obtained in two of the eight rounds. The distributions were skew-normal with parameters (ξ, ω, α) appropriately chosen so that the expected value would be equal to 6%, the standard deviation would be equal to

20%, and the skewness coefficient equal to some target level in the [-1;1] interval.

We chose a graphical representation of outcomes, with the possibility to initially sample from the displayed distributions, to make information provision easier to process for participants relative to a static numerical representation (on this see Kaufmann et al. (2013) and Bradbury et al. (2015)). Furthermore, we decided to present the probability density functions instead of histograms (like Holzmeister et al. (2020)) because their smoothness improves the comparability between different distributions.

Like in Brünner et al. (2011), higher skewness implies third-degree stochastic dominance for a given mean and variance of the distribution. Thus, for any decision maker with utility function $U(w)$ such that $U^{'}(w) > 0$, $U^{''}(w) < 0$, and $U^{'''}(w) > 0$, the distribution with higher skewness coefficient should be preferred (Levy, 1992). Menezes et al. (1980) define downside risk as *"a leftward transfer of risk, keeping mean and variance the same"*. Given two distributions with densities f(x) and g(x), if f dominates g by third-degree stochastic dominance (TSD), and they have the same mean and variance, then f has less downside risk than g, and it is more right-skewed than g (Menezes et al., 1980). Thus, the downside risk is decreasing in skewness in our framework, and a sufficient condition to prefer the distribution with a higher skewness coefficient is $U^{'''}(w) > 0$, or downside risk aversion (Menezes et al., 1980, theorem 2).

Table 1 reports a brief description of the rounds and their names. For a full description of the rounds, refer to the appendix.

In the rounds *Binary-base* and *Binary-adjustment*, subjects faced a binary decision between two distributions differing in skewness: they visualized both distributions (like those in Figure 1) and clicked on a button to choose the distribution they preferred. In the rounds *Multiple-base*, *Multiple-partial*, and *Multiple-full*, subjects made a decision among eleven distributions differing in skewness: they could visualize one distribution at a time and change the displayed distribution by horizontally dragging the slider placed below the distribution (see Figure 2).

In rounds *Skew-risk* and *Skew-risk-reference*, subjects could move two sliders to change the skewness and standard deviation levels. The skewness slider worked like in the previous rounds, while the standard deviation slider was vertical (Figure 3).

Round	Name	Description
1	*Binary-base*	Binary decision between 2 distributions differing in skewness
2/3	*Binary-adjustment*	Binary decision between 2 distributions differing in skewness and expected return
4	*Multiple-base*	Decision among 11 distributions differing in skewness
5	*Multiple-partial*	Decision among 11 distributions differing in skewness with one piece of information
6	*Multiple-full*	Decision among 11 distributions differing in skewness with three pieces of information
7	*Skew-risk*	Decision among 15 distributions differing in skewness and standard deviation
8	*Skew-risk-reference*	Decision among 15 distributions differing in skewness and standard deviation, with the provision of a reference point

Table 1: Description of the rounds

3.3 Description of the sample

We collected 180 valid observations in the month of February 2022 on the platform Prolific[5].

In the previous paragraphs, we outlined some issues stemming from using continuous distributions to represent alternatives, which may lead to a negative *"global preference"* for positively skewed distributions. Ganzach (2000) found that when individuals evaluate unfamiliar assets, the global preference affects both risk perception and perceived return. Since this does not happen in evaluations of familiar assets, it is important that the decision-makers have some familiarity with the decisional environment. Thus, we set some restrictions from the Prolific subject pool to account for the complexity of the task. Eligible subjects needed to be at least 21 years old, be fluent in English, have completed high school, and specialize in either a STEM subject or economics/finance. Moreover, the sample was gender-balanced. On average, subjects took 22 minutes (with a sd of 10 minutes) to complete the experiment and earned about £4.6. The final payment was composed of a fixed participation fee, a variable fee for part one (mean £1.1 and sd £0.31), and a variable fee for part two (mean £0.78 and sd £0.38).

[5]An observation is considered valid if the subject completed the study. This required passing the comprehension check. Forty-nine subjects did not pass the comprehension check.

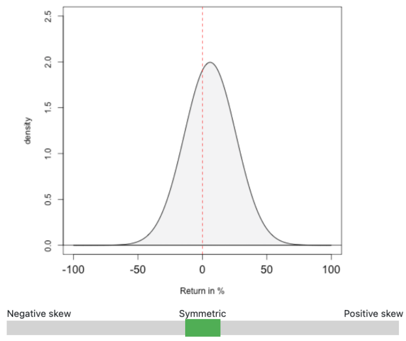

Figure 2: Decision framework of the *Multiple-base* round. The slider is moved to the left to reduce the skewness coefficient of the distribution or to the right to increase it. The range of the skewness coefficient is [-0.95, +0.95].

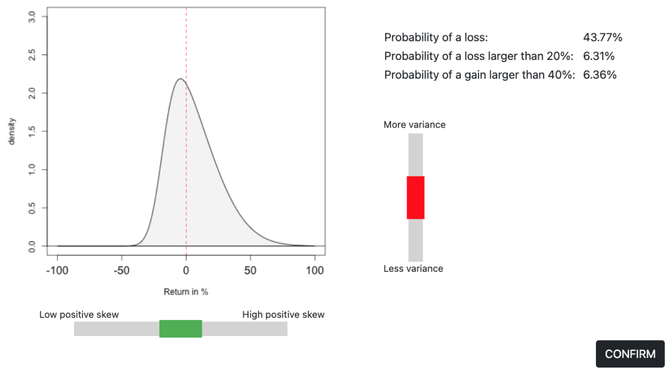

Figure 3: Decision framework of the *Skew-risk* (SR) and *Skew-risk-reference* (SRR) rounds. The horizontal slider is used to change the skewness coefficient, and the vertical slider is used to change the standard deviation. The probability to experience a loss, a loss larger than 20% and and gain larger than 40% update accordingly.

4 Results

The experiment is designed so that the complexity of the decision framework increases throughout the rounds. Since these final rounds offer the most insightful results of this research, we will start commenting on them from the last rounds and move backward.

4.1 Skewness and risk-taking

In the rounds *Skew-risk* (SR) and *Skew-risk-reference* (SRR), subjects could select both skewness and standard deviation. In one round, they were assigned to the positive skewness treatment, meaning they could choose only a positively skewed distribution, while in the other round, they were assigned to the negative skewness treatment, meaning they could choose only a negatively skewed distribution. The order was randomized at the subject level. For both treatments, the subjects could choose three levels of standard deviation: 0.16, 0.20, or 0.24, and five levels of skewness[6], for a total of fifteen possible distributions. Moreover, in the *Skew-risk-reference* (SRR) round, subjects were also provided with a reference point: their current return[7].

We found that, in both rounds, risk-taking (i.e., the level of chosen standard deviation) was significantly higher for the subjects assigned to the negative treatment (p-value of Wilcoxon rank sum test with continuity correction is 0.002^{**} in the SR round, and 0.018^{*} in the SRR round).

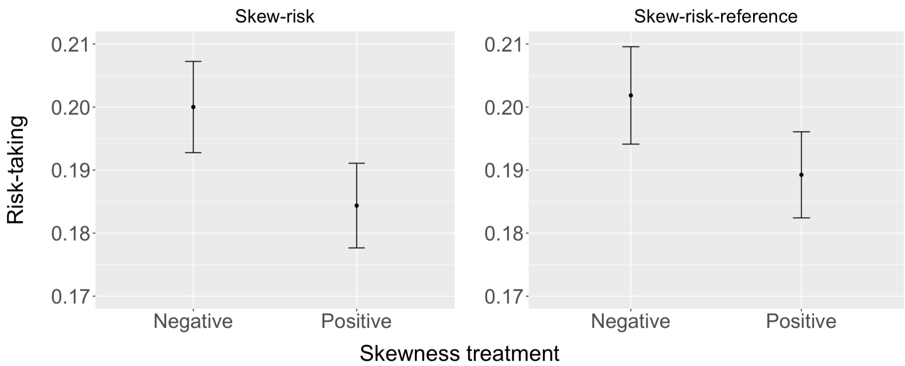

Figure 4: Average standard deviation, distinguishing by round and skewness environment. Subjects assigned to the negative treatment took on more risk.

Furthermore, we analyzed the level of skewness and standard deviation within treatment: although the level of these two variables was positively correlated in both rounds (controlling for treatment), the Kendall's Tau correlation coefficient was significantly different from zero only in the *Skew-risk-reference* round: $\tau = 0.13$ in round SR-positive, $\tau = 0.11$ in round SR-negative, $\tau = 0.30^{***}$ in round SRR-positive, $\tau = 0.34^{***}$ in round SRR-negative[8].

[6] The skewness coefficients of the distributions were 0.3, 0.5, 0.75, 0.85, and 0.95 for the positive treatment and -0.3, -0.5, -0.75, -0.85, and -0.95 for the negative treatment.

[7] The current return was a random draw from a distribution they had selected in one of the previous seven rounds. They were told their payoff of part one of the experiment would have been equal to the current return plus the realized return of the *Skew-risk-reference* round.

[8] In the Appendix we argue that there is stronger evidence of a positive correlation between

4.2 Analysis of skewness preferences

In the rounds *Binary-base* and *Multiple-base*, subjects were proposed investment opportunities, represented by a picture of their probability density functions. In the *Binary-base* round, the choice was between two alternatives: depending on the randomly assigned treatment, the choice could be between (i) a positively skewed and a symmetric distribution; (ii) a negatively skewed and a symmetric distribution, or (iii) a positively and negatively skewed distribution. All investments had the same expected return and volatility but differed in skewness (equal to -0.75, 0, and 0.75, depending on the distribution). Subjects could generate random samples from the displayed distributions to enhance familiarity with probability density functions. In the *Multiple-base* round, they could choose among eleven alternatives: five negatively skewed, one symmetric, and five positively skewed[9].

Since the distributions respected the skewness comparability criteria, an individual with $U^{'''}(w) > 0$ should have chosen the distribution with the largest skewness coefficient. In the *Binary-base* round, the proportion of subjects picking the most skewed distribution in the binary choice was not significantly different from 50%, neither at the aggregate level nor dividing by the three treatments. Therefore, we did not find evidence of the prevalence of prudence or skewness-seeking behavior (in the positive versus negative treatment). Moreover, we cannot reject the hypothesis that choices were made randomly in this round.

On the contrary, in the *Multiple-base* round, we found evidence of skewness-seeking behavior: the proportion of subjects investing in a positively skewed alternative was significantly different from 50% and equal to about 66%. Moreover, we can reject the hypothesis that choices were made randomly (p-value of Chi-squared test $< 0.001^{***}$).

4.3 Skewness preferences and information set manipulations

In the *Multiple-base* round, subjects could decide among eleven investment opportunities differing in skewness. Still, they had no numerical information about them,

---

skewness and risk-taking than the Kendall's τ suggests, even in the SR round.

[9]Expected return and variance were constant across alternatives, while skewness differed across alternatives. Skewness coefficients were -0.95, -0.85, -0.75, -0.50, -0.30, 0, 0.3, 0.5, 0.75, 0.85, 0.95. Kurtosis was the same for each couple of distributions with the same absolute level of skewness.

except that they all had the same expected return and volatility. All they could do was infer some information from the probability density functions. In the *Multiple-partial* round, subjects faced the same decision environment but received one piece of information regarding the displayed distributions. Depending on the treatment, they may visualize (i) the probability of a loss, (ii) the probability of a loss larger than 20% ("large loss"), or (iii) the probability of a gain larger than 40% ("large gain"). As they moved the slider, the distribution displayed changed, as well as the displayed probability and the area associated with the displayed probability (see Figure 5).

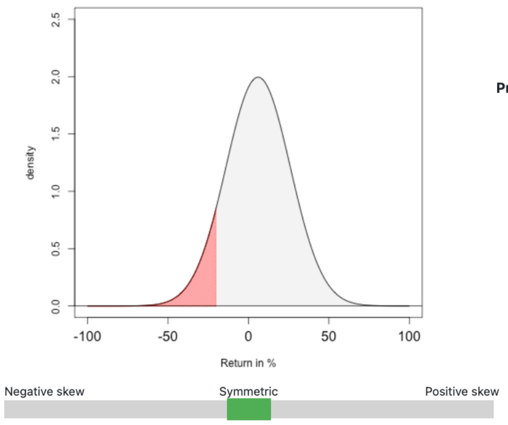

Figure 5: Decision framework of the *Multiple-partial* round. One piece of information is displayed.

First, we reject the hypothesis that choices were made randomly (p-value of Chi-squared test $< 0.001^{***}$). If subjects could perfectly (and somewhat unrealistically) infer probabilities from the pictures of the distributions, then none of the treatments should impact their decisions, which should be equal to the decision of the *Multiple-base*. On the contrary, if subjects had no understanding of the distributions, they would rely just on the probabilities shown and pick an edge choice - either the most positively or the most negatively skewed alternative - optimizing for the probabilities displayed. Assuming subjects fall somewhere in between, that is, they have some understanding of the probability density functions and they incorporate the new piece of information, we should expect the "probability of a loss" treatment to reduce skewness-seeking. In contrast, the other two treatments should increase it. Statistical tests do indeed indicate that treatments were effective in orienting decisions. Since the median skewness level across treatments is statistically different (p-value Kruskal-Wallis test $\leq 0.001^{***}$), we perform a pairwise comparison of the three treatments using a Wilcoxon rank sum test. The p-values $\{< 0.001^{***}, < 0.001^{***}, 0.23\}$ indicate that the difference in median skewness is insignificant only between the "probability of a large gain" and "probability of a large loss". Two-sample Kolmogorov-Smirnov tests on the distribution of skewness choices across the treatments confirm the previous results: p-values $= \{< 0.001^{***}, < 0.001^{***}, 0.42\}$. Despite the effectiveness of the treatments in orienting decisions, the choices at the *Multiple-partial* round were still consistent with those made at *Multiple-base* round: for all three treatments, subjects previously classified as skewness-seekers were still more skewness-seekers than the others (p-value of Wilcoxon test $< 0.01^{**}$ for all three treatments).

In the *Multiple-full* round, subjects were again asked to choose one of the eleven distributions, but they were provided with all three pieces of information about the treatments of the *Multiple-partial* round. This round was more complex than the previous because the three probabilities displayed provided contrasting cues: the probability of a large gain and the probability of a loss were increasing in skewness, while the probability of a large loss was decreasing in skewness. The increase in the complexity was indeed perceived by the subjects, who moved the slider (i.e., explored the environment) significantly more times. We can reject the hypothesis that decisions were affected by the treatment assigned at the *Multiple-partial* round (p-value of Kruskal-Wallis rank sum test equal to 0.63) and that the choices were random (p-value Chi-squared test $< 0.014^{*}$). This suggests that subjects incorporated new information when making their decision. We found that even if the skewness-seekers group was still the largest, the proportion of skewness-seekers reduced significantly from the *Multiple-base* to the *Multiple-full* round. The prevalence of skewness-seeking behavior largely disappeared: skewness-seekers were not significantly more than 50% of the sample anymore; they were about 46% versus 43% skewness-avoiders and 11% skewness neutral (i.e., those who chose the symmetric option). Subjects classified as skewness-seekers and skewness-avoiders based on the *Multiple-base* round were still on average skewness-seekers and skewness-avoiders, respectively, but the median skewness chosen for both groups was closer to zero. This process resulted in a more uniform (but not random) distribution of choices (Figure 6).

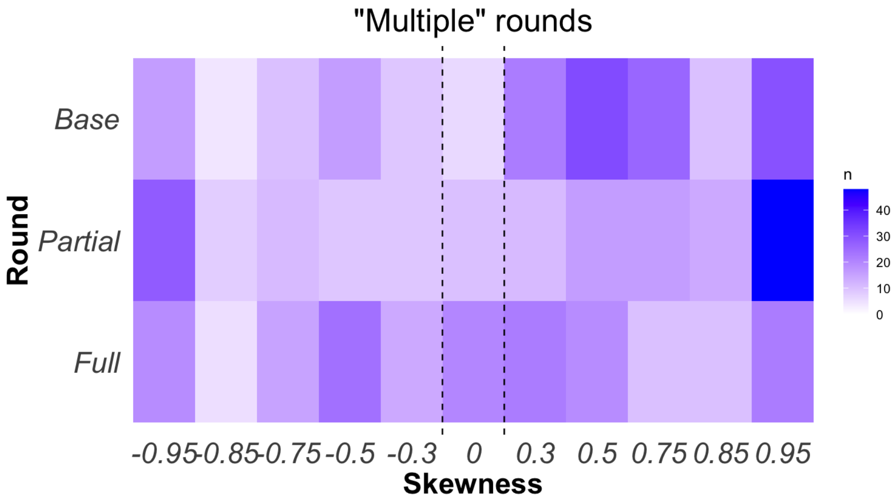

Figure 6: Heatmap of the choices at rounds *Multiple-base*, *Multiple-partial* and *Multiple-full*.

4.4 Trade-off between skewness and expected return

In the two *Binary-adjustment* rounds, subjects faced the same decision environment of the *Binary-base* round. In one round they could receive a bonus if they decided to invest in the opportunity they had not selected in the *Binary-base* round ("bonus treatment"). In contrast, in the other, they would pay a penalty if they decided to invest in the opportunity they had selected in the *Binary-base* round ("penalty treatment"). Both treatments were played during the two consecutive rounds, with the order of treatments being randomly assigned at the subject level. The value of the adjustment (bonus/penalty) ranged between 0% and 1%, extremes included, with 0.10% increments. In the first three columns of Table 2, we report the generalized linear mixed-effect models analyzing the probability of changing distribution with respect to the *Binary-base* round. The results show that subjects traded-off skewness with expected returns: the probability of change increases in the magnitude of the adjustment. Moreover, subjects were less likely to change when the shift would have been from a positively to a negatively skewed distribution (or vice versa) than when the change involved a symmetric distribution: the further away the new level of skewness from the initial favorite level, the less likely the subjects to change, given the adjustment.

4.5 Determinants of skewness preferences

We now combine decisions taken in multiple rounds to study the determinants of skewness preferences. Considering the decisions in the three *Binary* rounds (see Table 2, Models 4 and 5), we find that in general the higher the upside score, the more likely a subject to pick the most skewed distribution in a binary choice. The upside score is an indicator of the willingness to achieve better upside opportunities or speculative gains, which is computed based on the answers given in the final questionnaire.

Furthermore, we pool together the decisions made in the three *Multiple* rounds (Table 3). The upside score was a driver of skewness preferences also in these rounds. Moreover, risk perception of positive skewness was a driver of decisions: the higher the risk perception (of positively skewed distributions vis-a-vis negatively skewed distributions), the lower the skewness level and the lower the likelihood of choosing a positively skewed distribution. In the rounds *Multiple-partial* and *Multiple-full*, subjects were given different information about the distributions. We conjecture that subjects gave different weights to the three probabilities and, depending on these weightings, made a decision in the *Multiple-full* round. We use information collected in the aspirations phase to classify subjects into three groups, with each group expected to give more weight to one of the three probabilities. On the aspirations page, we asked subjects whether they would rather combine ownership of a stock with a financial instrument that enhances returns in case the stock performs well (i.e., a call option) or with a financial instrument that reduces losses in case the stock performs badly (i.e., a put option). We assume that those who chose the call option are relatively more focused on the upside, while those who chose the put option are relatively more focused on the downside. Therefore, more upside-focused subjects are expected to give more weight to the probability of a large gain, so they are expected to be relatively more skewness-seeker. As for the downside-

	Model 1	Model 2	Model 3	Model 4	Model 5
Intercept	-1.64^{*}	0.13	-2.61^{***}	-1.24	-0.27
	(0.79)	(1.42)	(0.78)	(1.25)	(0.55)
Adjustment	2.41^{***}	2.38^{***}	2.44^{***}		
	(0.64)	(0.64)	(0.64)		
Upside	0.49^{*}	0.49^{*}	0.55^{*}	0.44^{*}	0.08^{*}
	(0.23)	(0.23)	(0.24)	(0.20)	(0.04)
Round 3	1.02^{***}	1.02^{***}	1.01^{***}		
	(0.30)	(0.30)	(0.30)		
Treatment penalty	0.67^{*}	0.67^{*}	0.68^{*}		
	(0.29)	(0.29)	(0.29)		
Treat Pos and Neg	-0.98^{*}	-0.97^{*}		-0.21	0.41^{***}
	(0.45)	(0.44)		(0.40)	(0.07)
Treat Pos and Symm	-0.40	-0.38		-0.49	0.67^{***}
	(0.43)	(0.43)		(0.38)	(0.07)
Skewness-seeker			-0.40		
			(0.37)		
Return most skewed					0.42^{***}
					(0.07)
Return least skewed					-0.47^{***}
					(0.07)
Perceived gap				0.05^{*}	
				(0.02)	
Controls	Yes	Yes	Yes	Yes	Yes
AIC	442.26	440.21	449.56	251.22	836.30
BIC	485.00	482.96	476.76	279.95	887.80
Log Likelihood	-210.13	-209.11	-217.78	-116.61	-406.15
Num. obs.	360	360	360	180	540
Num. groups: participant	180	180	180		180
Var: participant (Int)	2.02	1.98	2.34		0.08
Deviance				233.22	
Var: Residual					0.19

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 2: In Models 1-3 (GLMM), we study the probability to change distribution in the *Binary-adjustment* rounds with respect to the *Binary-base* round. In Model 4 (GLM), we model the probability to choose the distribution with the largest skewness coefficient in the *Binary-base*, and in Model 5 (LMM) we model the skewness coefficient of the chosen distribution in the *Binary-base* and *Binary-adjustment* rounds.

focused subject, we distinguish between subjects concerned about avoiding losses and subjects concerned about avoiding large losses. Hence, we consider the maximum self-reported threshold for losses in an investment: those who are willing to bear a loss up to a given low threshold τ are expected to give more weight to the probability of a loss, so they are expected to be relatively less skewness-seeker. The others, who can tolerate losses larger than τ, are expected to attribute more weight to the probability of a large loss, so they are expected to be relatively more skewness-seeker. A natural threshold could be 10%, that is the mid-point between the two loss levels provided, 0% and 20%. Since the probability of a loss has been found to be salient in risky decisions (Holzmeister et al., 2020; Zeisberger, 2022), we also used a 5% threshold, which is closer to the salient value of 0%.

The regressions presented in Table 3 are consistent with our conjectures. Skewness preferences seem to be ultimately driven by two channels: first, the speculative channel (operationalized through the "Upside score" and the "Call preference" dummy) indicates that speculators tend to choose more positively skewed distributions than non-speculators. The second channel is risk perception of positively skewed distributions versus negatively skewed distributions, a variable that we call "Risk perception". While risk perception plays an important role in skewness preferences, risk preferences do not seem to play a role within our framework. In part two of the experiment, we elicited risk preferences using a modified version of the multiple price list (Holt and Laury, 2002), and classified subjects as risk-averters, risk-neutral, and risk-seekers. Across the rounds, we found no systematic difference in skewness preferences between the three groups. The result is not surprising since skewness concerns downside risk preference, and both risk lovers and risk averters can be downside risk averse (Menezes et al., 1980). Indeed Haering et al. (2020) did not find differences across these two groups in their preferences for higher-order odd moments, including skewness.

	Model 1	Model 2	Model 3	Model 4
Intercept	0.33	0.60^{*}	0.27	0.53^{*}
	(0.28)	(0.26)	(0.28)	(0.27)
Risk perception	-0.15^{***}	-0.15^{***}	-0.13^{***}	-0.15^{***}
	(0.03)	(0.03)	(0.03)	(0.03)
Upside	0.11^{**}			
	(0.04)			
Call pref		0.15^{*}	0.41^{***}	0.20^{*}
		(0.07)	(0.12)	(0.09)
Put pref and high loss (5%)			0.32^{**}	
			(0.12)	
Put pref and high loss (10%)				0.09
				(0.10)
No information shown	0.21^{***}	0.21^{***}	0.21^{***}	0.21^{***}
	(0.05)	(0.05)	(0.05)	(0.05)
P large gain shown	0.48^{***}	0.48^{***}	0.49^{***}	0.48^{***}
	(0.08)	(0.08)	(0.08)	(0.08)
P large loss shown	0.57^{***}	0.56^{***}	0.56^{***}	0.56^{***}
	(0.08)	(0.08)	(0.08)	(0.08)
P loss shown	-0.52^{***}	-0.52^{***}	-0.52^{***}	-0.52^{***}
	(0.08)	(0.08)	(0.08)	(0.08)
Controls	Yes	Yes	Yes	Yes
AIC	1000.61	1002.31	999.81	1006.23
BIC	1056.40	1058.10	1059.89	1066.31
Log Likelihood	-487.30	-488.15	-485.90	-489.12
Num. obs.	540	540	540	540
Num. groups: participant	180	180	180	180
Var: participant (Int)	0.11	0.11	0.11	0.11
Var: Residual	0.25	0.25	0.25	0.25

$^{***}p < 0.001$; $^{**}p < 0.01$; $^{*}p < 0.05$

Table 3: Models of the level of skewness chosen at rounds *Multiple-base*, *Multiple-partial*, and *Multiple-full*. The four specifications test the speculative channel in different ways.

5 Discussion

5.1 Skewness preferences

We defined skewness-seeking behavior as the preference for a positively skewed distribution over a symmetric and a negative one when multiple alternatives with the same mean and variance were available. Our definition is still consistent with Ebert and Wiesen (2011): when a subject picked a positively skewed distribution in the *Multiple* rounds, she was not picking another distribution with the same mean, variance, and kurtosis. In the *Binary-base* round, subjects faced a binary decision, and skewness-seeking could only be tested for the treatment where a positively and a negatively skewed distribution were available. In the other two treatments, only prudence could be tested. In the *Binary-base* round, in none of the three treatments, we found evidence of subjects choosing the distribution with the largest skewness coefficient with a probability significantly larger than 50%. Therefore, we rejected the prevalence of both prudence and skewness-seeking behavior. However, a few considerations are worth mentioning. Firstly, our results already contrast Vrecko et al. (2009) and Holzmeister et al. (2020), who found evidence of skewness-avoidance when probability density functions are used. Secondly, the preference for skewness may not arise immediately: in Brünner et al. (2011) the proportion of skewness-seekers in the first two rounds was about 40%, while in the last rounds, it was around 67%, and since the order of the rounds was randomized, they attributed this phenomenon to some form of learning. Finally, since Ebert and Wiesen (2011) suggest that prudent individuals are mostly skewness-seekers, but the opposite is not necessarily true, we could see the proportion of prudent individuals in the *Binary-base* round as a lower bound for skewness-seekers since the imprudent skewness-seekers should be more than the prudent skewness-avoiders. Indeed, when we tested for skewness-seeking behavior in the *Multiple-base*, 80% of prudent decision-makers (i.e., those who made a prudent choice in the *Binary-base* round) were skewness-seekers, while 55% of imprudent were skewness-seekers. Overall, in the *Multiple-base* round, about two-thirds of the subjects were skewness-seekers, a proportion in line with the existing experimental literature. Considering the five pairs of distributions with the same mean, variance, and kurtosis, the proportion of skewness-seekers in each pair ranged between 65% and 72%. The prevalence of skewness-seeking behavior can be attributed to at least two elements. First, compared to the *Binary-base* round, subjects earned some additional experience with distributions differing in skewness, and the round properly allowed the testing of this behavior. Second, in the *Multiple-base* round subjects could visualize only one distribution at a time, and by dragging the slider they could change the displayed distribution: the impact of increasing skewness on the probability density function became extremely evident and salient. Thus, the "dynamic presentation" of the alternatives may have made the comparisons easier compared to the traditional "static presentation".

If subjects had mean-variance preferences, then they should have always changed distribution from the *Binary-base* to the two *Binary-adjustment* rounds because they could have obtained an investment with a higher expected return and the same volatility. Nevertheless, the results show that subjects did not necessarily pick the distribution with the higher expected return. The probability of changing distribution was increasing in the magnitude of the adjustment, indicating that subjects traded-off skewness for expected return. Thus, skewness matters in subjects' decisions, and a utility function incorporating only mean and variance would be inconsistent with our results.

Throughout the rounds, subjects exhibited consistency in their revealed preferences: we conduct seemingly unrelated regressions (SUR) of the skewness levels of the six rounds in which the distributions had the same expected return (i.e., all the rounds except the two *Binary-adjustment* rounds), and analyze the correlation matrix of the residuals. The residuals are all positively correlated. This correlation is statistically significant in most cases, indicating that unobservable idiosyncratic traits influenced the decisions across the rounds (see the appendix for more details).

5.2 Risk perception

Risk perception played a relevant role in explaining skewness preferences. Within our framework, the relationship between skewness and risk would depend on the operationalization of the latter. Volatility σ was constant across the distributions, and the behavioral risk measure σ_B^2 (Davies and De Servigny, 2012) was decreasing in skewness, as well as semivariance. The probability to experience a loss -

which, according to Holzmeister et al. (2020), is the main channel in which skewness translates into risk perception - was increasing in skewness. Finally, since Oja's skewness comparability criterion was satisfied, a distribution with a higher skewness coefficient could be considered a downside risk decrease with respect to any other distribution with a lower skewness coefficient.

In the *Multiple* rounds, risk perception was correlated with actual choices. In particular, in the *Multiple-full*, i.e., when information became "fully available", the median skewness level for subjects who perceived positive skewness as less risky was significantly higher than zero ($p < 0.003^{**}$), while it was significantly lower than zero for subjects who perceived positive skewness riskier ($p < 0.032^{*}$). Only about 35% perceived positively skewed distributions as riskier than negatively skewed distributions, while 26% perceived them as safer, and the remaining 39% of the subjects believed they bear about the same risk. This result is in line with the idea that risk can be measured in different ways, but it contrasts with previous literature specifically focused on risk perception and skewness of continuous distributions (Holzmeister et al., 2020). This difference could stem from the elicitation procedure: while the measurement of risk perception was a key element of Holzmeister et al.'s study, which was elicited for every distribution, we only asked our subjects at the end of the study to indicate their level of agreement with the statement that positively skewed distributions are riskier than negatively skewed ones. Our subjects had thus already experienced the treatments and were aware of the relationship between skewness and the probability of losses, large losses, and large gains. However, at least part of the risk perception was already formed before the visualization of information about probabilities, as risk perception was significantly correlated with choices made before probabilities were provided (*Multiple-base* round). Like Holzmeister et al. (2020), we found that investment propensity and risk perception are inversely related.

5.3 Skewness and risk-taking: substitutes or complements

Most experimental evidence indicates a positive relationship between skewness and risk-taking (Åstebro et al., 2015; Brünner et al., 2011; Dertwinkel-Kalt and Köster, 2020; Ebert and Wiesen, 2011; Ebert, 2015). Grossman and Eckel (2015) also found that risk-taking increases as skewness available increases, but Taylor (2020) while confirming this result, attributed part of this effect to loss aversion.

The justification of the positive relationship between skewness and risk-taking can be found in Amaya et al. (2015): *"As positive asymmetry increases, volatility is welfare increasing as it implies a larger probability of an extremely good state of the economy. The opposite is true for the case of negative skewness since higher volatility increases the likelihood of a left tail event"*. Behavioral components may enhance this mechanism: according to Åstebro et al. (2015), skewness-seeking behavior is driven by optimism and likelihood insensitivity. Similarly, Dertwinkel-Kalt et al. (2020) relate skewness-seeking behavior to salience theory (Bordalo et al., 2012). On the contrary, Bougherara et al. (2021) found that variance and skewness do not interact in a positively skewed environment, while the difference between the certain equivalent of a highly negatively skewed prospect and a low negatively skewed prospect is higher when variance is higher. Bougherara et al. (2022) did not find a significant interaction between skewness and risk-taking, instead.

We contribute to this literature with two findings. First, we found that subjects took more risk when forced (by the treatment) to choose negatively skewed distributions. Such distributions have a longer left tail and a shorter right tail, so the probability of obtaining large positive outcomes is relatively lower. A subject who is interested in improving her upside may be forced to increase risk-taking because changing only skewness may not be satisfactory enough. This is not necessary in a positively skewed environment, where increasing skewness may be sufficient. Since a satisfactory upside may be achieved either by increasing the skewness coefficient in a positively skewed environment or by increasing variance in a negatively skewed environment, we claim that skewness-direction and risk-taking can be seen as substitutes.

Secondly, we have found that when a reference point is provided, skewness and risk-taking are positively correlated: individuals react to the reference point by picking a corner distribution more often. We define a corner distribution as a distribution where both skewness and standard deviation are either maximized or minimized within the available set of distributions. In this sense, skewness and risk-taking acted as complements to achieve a goal: in the maximization case, subjects maximized the probability of a large gain, whereas in the minimization case, they minimized the probability of a loss[10]. For each heatmap in Figure 7, the top-right and bottom-left squares represent the corner distributions: for both treatments, they were more frequent in the *Skew-risk-reference* rounds (heatmaps on the right) than in the *Skew-risk* rounds (heatmaps on the left).

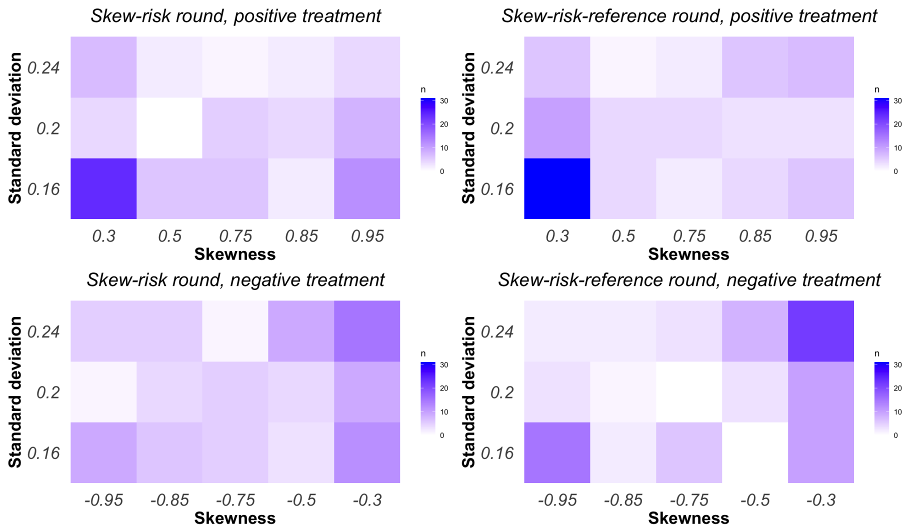

Figure 7: Heatmaps of the skewness and standard deviation levels chosen in the *Skew-risk* and *Skew-risk-reference* rounds, distinguishing for assigned skewness treatment.

6 Conclusions

Our experiment investigated preferences for continuous distributions differing in skewness. Our results are manifold, and they concern two main areas: the analysis of skewness preferences, and the interactions between skewness and risk-taking.

In our main settings - choice across multiple distributions based solely on the plots of the probability density functions - we found the prevalence of skewness-seeking behavior, with a proportion of skewness-seekers consistent with the pre-

[10]We did not find a specific relationship between the level of the reference point and the choices of skewness and standard deviation. Indeed, several subjects would have maximized/minimized both anyway (like they did in the *Skew-risk* round when there was no feedback), and there is not a unanimous theoretical relationship between the level of the reference point and the reaction: a large current return may indeed be consistent both with the house money effect (Thaler and Johnson, 1990), which would suggest an increase in risk-taking, as well as with prospect-theory-like preferences (Tversky and Kahneman, 1992), which would suggest a decrease in risk-taking.

vious literature. However, unlike previous contributions, we found this behavior using continuous distributions defined on a choice support that included gains and losses and giving the possibility to choose positively skewed, symmetric, and negatively skewed distributions. Consistent with the literature, our results suggest that a positively skewed distribution may be more attractive than a negatively skewed distribution because of its longer right tail, which represents the speculative channel, and its shorter left tail, which may represent the risk channel. Indeed, we show that several subjects perceive positively skewed distributions as safer than negatively skewed ones. Skewness preferences would therefore represent a kind of meta-preferences summarizing the view of the subject of the *skewness trade-off*. In other words, the choices represent each subject's synthesis of the features inferred from the probability density functions and the probabilities provided. This synthesis is the result of a weighting process of the (un)desirable characteristics. Quantification of the weights as well as the identification of other sources is beyond the scope of the paper, and it is left for future research.

Finally, we found a twofold relationship between skewness and risk-taking. First, the environment in which decisions are made significantly affects risk-taking: individuals are more risk-taking in a negatively skewed environment than in a positively skewed one. This finding, in sharp contrast to the existing literature, is related to the idea that since negatively skewed distributions have a relatively short right tail, the only way to increase the probability of large gains is to resort to more risk-taking, thus increasing the dispersion parameter. Thus, volatility is a kind of substitute for positive skewness. Moreover, when subjects are provided with a reference point, the choices of skewness and standard deviation tend to be more positively correlated, jointly working to either maximize the probability of large outcomes or minimize the probability of a loss. This last piece of evidence suggests a potential connection between Prospect Theory (Tversky and Kahneman, 1992) and skewness. An investigation of this connection goes beyond the scope of this paper, but further research in this direction is needed.

References

D. Amaya, P. Christoffersen, K. Jacobs, and A. Vasquez. Does realized skewness predict the cross-section of equity returns? *Journal of Financial Economics*, 118 (1):135–167, 2015. ISSN 0304-405X. Publisher: Elsevier.

F. D. Arditti. Risk and the required return on equity. *The Journal of Finance*, 22 (1):19–36, 1967. ISSN 0022-1082. Publisher: JSTOR.

N. Barberis and M. Huang. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5):2066–2100, 2008. ISSN 0002-8282. Publisher: American Economic Association.

N. C. Barberis. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1):173–196, 2013. ISSN 0895-3309. Publisher: American Economic Association.

P. Bordalo, N. Gennaioli, and A. Shleifer. Salience theory of choice under risk. *The Quarterly journal of economics*, 127(3):1243–1285, 2012. ISSN 0033-5533. Publisher: MIT Press.

D. Bougherara, L. Friesen, and C. Nauges. Risk taking with left-and right-skewed lotteries. *Journal of Risk and Uncertainty*, 62(1):89–112, 2021. ISSN 0895-5646. Publisher: Springer.

D. Bougherara, L. Friesen, and C. Nauges. Risk-taking and skewness-seeking behavior in a demographically diverse population. *Journal of Economic Behavior & Organization*, 201:83–104, 2022. ISSN 0167-2681. Publisher: Elsevier.

B. Boyer, T. Mitton, and K. Vorkink. Expected idiosyncratic skewness. *The Review of Financial Studies*, 23(1):169–202, 2010. ISSN 0893-9454. Publisher: Society for Financial Studies.

B. H. Boyer and K. Vorkink. Stock options as lotteries. *The Journal of Finance*, 69(4):1485–1527, 2014. ISSN 0022-1082. Publisher: Wiley Online Library.

M. A. Bradbury, T. Hens, and S. Zeisberger. Improving investment decisions with simulated experience. *Review of Finance*, 19(3):1019–1052, 2015. ISSN 1573-692X. Publisher: Oxford University Press.

P. L. Brockett and J. R. Garven. A reexamination of the relationship between preferences and moment orderings by rational risk-averse investors. *The Geneva Papers on Risk and Insurance Theory*, 23:127–137, 1998. ISSN 0926-4957. Publisher: Springer.

P. L. Brockett and Y. Kahane. Risk, return, skewness and preference. *Management Science*, 38(6):851–866, 1992. ISSN 0025-1909. Publisher: INFORMS.

T. Brünner, R. Levínský, and J. Qiu. Preferences for skewness: evidence from a binary choice experiment. *The European Journal of Finance*, 17(7):525–538, 2011. ISSN 1351-847X. Publisher: Taylor & Francis.

D. L. Chen, M. Schonger, and C. Wickens. oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016. ISSN 2214-6350. Publisher: Elsevier.

W. H. Chiu. Skewness preference, risk taking and expected utility maximisation. *The Geneva Risk and Insurance Review*, 35:108–129, 2010. ISSN 1554-964X. Publisher: Springer.

G. B. Davies and A. De Servigny. *Behavioral investment management: An efficient alternative to modern portfolio theory.* McGraw Hill Professional, 2012. ISBN 0-07-174660-9.

C. Deck and H. Schlesinger. Exploring higher order risk effects. *The Review of Economic Studies*, 77(4):1403–1420, 2010. ISSN 1467-937X. Publisher: Wiley-Blackwell.

M. Dertwinkel-Kalt and M. Köster. Salience and skewness preferences. *Journal of the European Economic Association*, 18(5):2057–2107, 2020. ISSN 1542-4766. Publisher: Oxford University Press.

S. Ebert. On skewed risks in economic models and experiments. *Journal of Economic Behavior & Organization*, 112:85–97, 2015. ISSN 0167-2681. Publisher: Elsevier.

S. Ebert and D. Wiesen. Testing for prudence and skewness seeking. *Management Science*, 57(7):1334–1349, 2011. ISSN 0025-1909. Publisher: INFORMS.

C. C. Eckel and P. J. Grossman. Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and human behavior*, 23(4):281–295, 2002. ISSN 1090-5138. Publisher: Elsevier.

C. C. Eckel and P. J. Grossman. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior & Organization*, 68(1):1–17, 2008. ISSN 0167-2681. Publisher: Elsevier.

L. Eeckhoudt and H. Schlesinger. Putting risk in its proper place. *American Economic Review*, 96(1):280–289, 2006. ISSN 0002-8282. Publisher: American Economic Association.

K. Fairley and A. G. Sanfey. The role of demographics on adolescents' preferences for risk, ambiguity, and prudence. *Journal of Economic Behavior & Organization*, 179:784–796, 2020. ISSN 0167-2681. Publisher: Elsevier.

Y. Ganzach. Judging risk and return of financial assets. *Organizational behavior and human decision processes*, 83(2):353–370, 2000. ISSN 0749-5978. Publisher: Elsevier.

T. A. Garrett and R. S. Sobel. Gamblers favor skewness, not risk: Further evidence from United States' lottery games. *Economics Letters*, 63(1):85–90, 1999. ISSN 0165-1765. Publisher: Elsevier.

J. Golec and M. Tamarkin. Bettors love skewness, not risk, at the horse track. *Journal of political economy*, 106(1):205–225, 1998. ISSN 0022-3808. Publisher: The University of Chicago Press.

P. J. Grossman and C. C. Eckel. Loving the long shot: Risk taking with skewed lotteries. *Journal of Risk and Uncertainty*, 51:195–217, 2015. ISSN 0895-5646. Publisher: Springer.

A. Haering, T. Heinrich, and T. Mayrhofer. Exploring the consistency of higher order risk preferences. *International Economic Review*, 61(1):283–320, 2020. ISSN 0020-6598. Publisher: Wiley Online Library.

T. Heinrich and J. Shachat. The development of risk aversion and prudence in Chinese children and adolescents. *Journal of Risk and Uncertainty*, 61:263–287, 2020. ISSN 0895-5646. Publisher: Springer.

C. A. Holt and S. K. Laury. Risk aversion and incentive effects. *American economic review*, 92(5):1644–1655, 2002. ISSN 0002-8282. Publisher: American Economic Association.

F. Holzmeister, J. Huber, M. Kirchler, F. Lindner, U. Weitzel, and S. Zeisberger. What drives risk perception? A global survey with financial professionals and laypeople. *Management Science*, 66(9):3977–4002, 2020. ISSN 0025-1909. Publisher: INFORMS.

C. Kaufmann, M. Weber, and E. Haisley. The role of experience sampling and graphical displays on one's investment risk appetite. *Management science*, 59(2): 323–340, 2013. ISSN 0025-1909. Publisher: INFORMS.

M. Kimball. Precautionary Saving in the Small and in the Large. *Econometrica*, 58:53–73, 1990.

H. Levy. Stochastic dominance and expected utility: Survey and analysis. *Management science*, 38(4):555–593, 1992. ISSN 0025-1909. Publisher: INFORMS.

J. C. Mao. Survey of capital budgeting: Theory and practice. *Journal of finance*, pages 349–360, 1970. ISSN 0022-1082. Publisher: JSTOR.

H. M. Markowitz. Foundations of portfolio theory. *The journal of finance*, 46(2): 469–477, 1991. ISSN 0022-1082. Publisher: JSTOR.

C. Menezes, C. Geiss, and J. Tressler. Increasing downside risk. *The American Economic Review*, 70(5):921–932, 1980. ISSN 0002-8282. Publisher: JSTOR.

C. N. Noussair, S. T. Trautmann, and G. Van de Kuilen. Higher order risk attitudes, demographics, and financial decisions. *Review of Economic Studies*, 81(1):325–355, 2014. ISSN 1467-937X. Publisher: Oxford University Press.

H. Oja. On location, scale, skewness and kurtosis of univariate distributions. *Scandinavian Journal of statistics*, pages 154–168, 1981. ISSN 0303-6898. Publisher: JSTOR.

J. W. Pratt. Risk aversion in the large and in the small. *Econometrica*, 32(1-2): 122–136, 1964.

C. Schneider and O. Spalt. Conglomerate investment, skewness, and the CEO long-shot bias. *The Journal of Finance*, 71(2):635–672, 2016. ISSN 0022-1082. Publisher: Wiley Online Library.

B. Summers and D. Duxbury. Peak impact: Financial risk perception and the peak of the return distribution. Technical report, Working Paper, 2006.

M. P. Taylor. Liking the long-shot... but just as a friend. *Journal of Risk and Uncertainty*, 61:245–261, 2020. ISSN 0895-5646. Publisher: Springer.

R. H. Thaler and E. J. Johnson. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management science*, 36(6):643–660, 1990. ISSN 0025-1909. Publisher: INFORMS.

A. Tversky and D. Kahneman. Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157):1124–1131, 1974. ISSN 0036-8075. Publisher: American association for the advancement of science.

A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5:297–323, 1992. ISSN 0895-5646. Publisher: Springer.

W. Van Zwet. Convex transformations of random variables, Math. *Centrum, Amsterdam*, 20, 1964.

D. Vrecko, A. Klos, and T. Langer. Impact of presentation format and self-reported risk aversion on revealed skewness preferences. *Decision Analysis*, 6(2):57–74, 2009. ISSN 1545-8490. Publisher: INFORMS.

S. Zeisberger. Do people care about loss probabilities? *Journal of Risk and Uncertainty*, 65(2):185–213, 2022. ISSN 0895-5646. Publisher: Springer.

T. Åstebro, J. Mata, and L. Santos-Pinto. Skewness seeking: risk loving, optimism or overweighting of small probabilities? *Theory and Decision*, 78:189–208, 2015. ISSN 0040-5833. Publisher: Springer.

Appendix

Appendix A - Experimental design

Introduction to the experiment

In the description of the study, we told subjects: *"In this study, we will ask you to make investment decisions. On top of the fixed payment, you will earn a bonus payment. The experiment will take between 20 and 25 minutes. We recommend the use of the Web Browser Chrome."*

As they clicked on the link we provided, they had to enter their Prolific ID, and after the 'Welcome page', in which they were briefly introduced to the content of the experiment, subjects had to complete an interactive tutorial, which included the meaning of a probability density function and the concepts of variance and skewness. Subjects could change these variables and learn the impact they had on the probability density function using two sliders like they would have done in some of the subsequent rounds.

Afterward, they were administered a comprehension check: unless they answered all questions correctly, they could not move forward. However, they could try to answer as many times as they wished (see Figure 8).

The tutorial served as a refresher of the previously acquired skills related to the interpretation of probability density functions. The comprehension check served the purpose to make sure that all subjects had a clear idea of what the distributions represented.

Aspirations

Before starting part one, subjects were asked about their aspirations for gains in a stock market investment, the maximum loss they would be willing to bear for a stock market investment, and whether they would rather combine the ownership of stocks with *"an option that increases your gains when the stock performs very well"* (i.e., a call option) or with *"an option that reduces your losses when the stock performs very bad"* (i.e., a put option).

Summary of the rounds

In each round, the subjects must choose one of the available distributions differing in skewness.

Binary-base (Round 1): binary choice between two distributions differing in skewness. There are three possible treatments that affect the skewness of the alternative: (-0.75, 0), (-0.75, +0.75), (0, +0.75)

Binary-adjustment (Rounds 2/3): binary choice between the two distributions of the *Binary-base* round. In one round the distribution previously chosen receives a random penalty, while in the other the distribution not chosen receives a random bonus. The order of the penalty/bonus and their magnitude are randomized. The adjustment ranges between 0% and 1%, with 0.10% increments.

Multiple-base (Round 4): multiple choice among eleven distributions differing in skewness with no additional information. Skewness levels are -0.95, -0.85, -0.75, -0.50, -0.30, 0, +0.30, +0.50, +0.75, +0.85, +0.95.

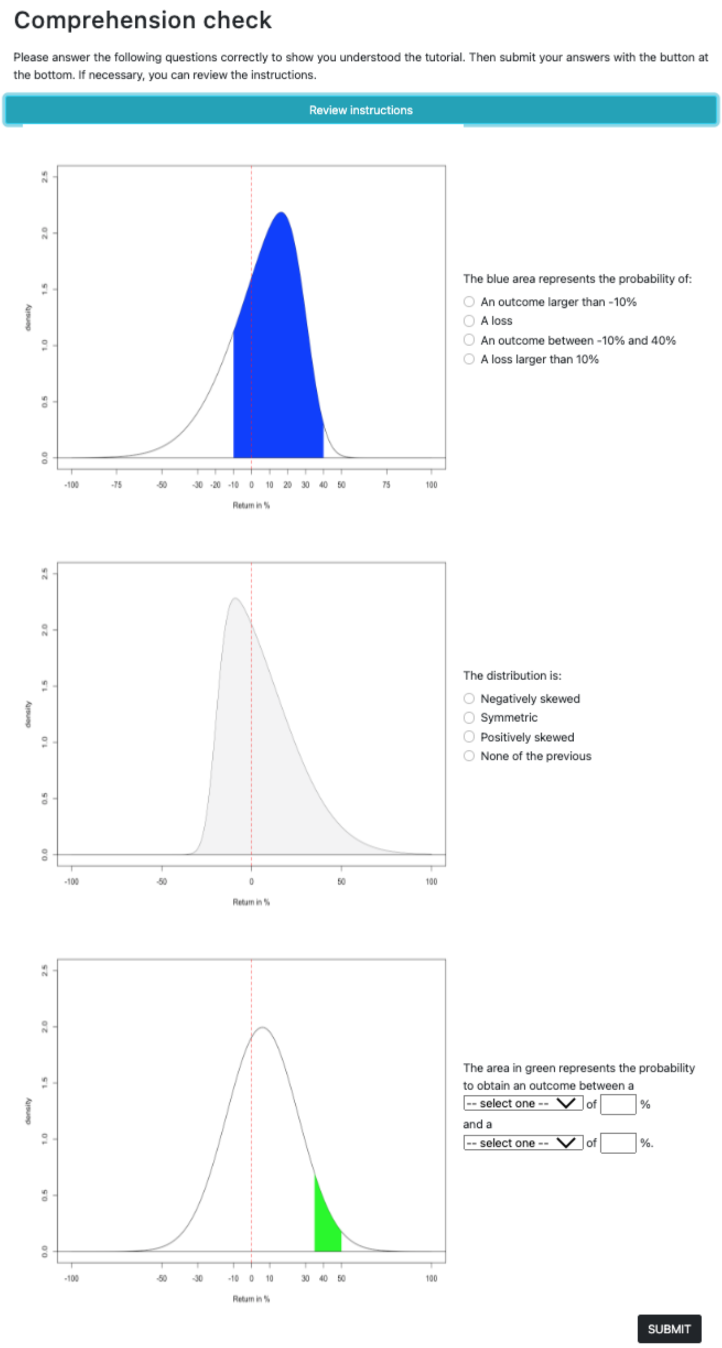

Figure 8: Comprehension check page

Multiple-partial (Round 5): same decision environment of the *Multiple-base* round, but one measure of risk/return shown. The three possible measures are the probability of a loss, the probability of a loss larger than 20% and the probability of a gain larger than 40%. The measure shown is chosen randomly for each subject.

Multiple-full (Round 6): same decision environment of the *Multiple-base* round, but all the three measures of risk/return previously indicated are shown.

Skew-risk (Round 7): multiple choice among fifteen distributions differing in skewness and standard deviation. The levels of standard deviations are 0.16, 0.20, 0.24, while the levels of skewness are -0.95, -0.85, -0.75, -0.50, -0.30 if the subject is assigned to the negative treatment, and +0.30, +0.50, +0.75, +0.85, +0.95 if she is assigned to the positive treatment. Treatment is assigned randomly at the subject level.

Skew-risk-reference (Round 8): same decision environment of the *Skew-risk* round, but the treatment assigned is now the opposite. Moreover, the current return accumulated is shown.

Sample generation process of the *Binary-base* round

In the *Binary-base* round (Round 1) subjects could generate samples from the two displayed distributions. The generation process of the samples worked in this way: ten random observations were drawn from a standard uniform distribution, they were ranked from the largest to the smallest, and then for each $u_{(10)}, u_{(9)}, ..., u_{(2)}, u_{(1)}$, $F^{-1}(u_{(j)})$ and $G^{-1}(u_{(j)})$ were shown to the subjects (where F and G are the cumulative distribution functions of the two displayed investment opportunities and $u_{(j)}$is the j^{th} order statistics). Subjects could generate as many samples as they wished.

Questionnaires

The subjects were asked to express their level of agreement with the statements reported in Figure 9. Investment-related questions should be answered hypothetically in case the subject did not have enough money to invest.

The upside score is computed based on answers to questions 1, 6, 7, and 9, and it indicates how much a subject focuses on the upside.

The downside score is computed based on answers to questions 3, 5, 10, and 12, and it indicates how much a subject focuses on the downside.

The soundness score is computed based on answers to questions 2, 4, 8 (reversed), and 11, and it indicates how reasonable a subject is about her approach to the financial markets.

The questionnaire above was followed by a demographic questionnaire asking about gender, age, race, profession, education level, education field and nationality.

Sample descriptive statistics

In Table 4 we summarize some of the characteristics of the subjects in our sample.

Statements	Disagree				Agree
When I invest, I care more about the possibility to obtain large gains than to suffer large losses	○ 1	○ 2	○ 3	○ 4	○ 5
I think a 7% average annual return for stocks in the long term is a very good result	○ 1	○ 2	○ 3	○ 4	○ 5
Avoiding large losses is my first target when I invest	○ 1	○ 2	○ 3	○ 4	○ 5
I consider my investment strategy prudent with respect to downside risks	○ 1	○ 2	○ 3	○ 4	○ 5
When I invest, I care more about avoiding losses than obtaining very large returns	○ 1	○ 2	○ 3	○ 4	○ 5
If I believe an asset is in a bubble, I buy that asset to make a profit before the bubble pops	○ 1	○ 2	○ 3	○ 4	○ 5
I think buying a lottery ticket is a good idea because I might become rich	○ 1	○ 2	○ 3	○ 4	○ 5
I am a gambler	○ 1	○ 2	○ 3	○ 4	○ 5
Cryptocurrencies represent more than 10% of my ideal portfolio	○ 1	○ 2	○ 3	○ 4	○ 5
I believe the ideal amount of cryptocurrencies to hold is less than 1% or even 0% of a portfolio	○ 1	○ 2	○ 3	○ 4	○ 5
I am willing to invest in a bond with a return lower than 1% or to hold a lot of cash to avoid investing too much in riskier assets, like stocks.	○ 1	○ 2	○ 3	○ 4	○ 5
I do not want to invest in the stock market because it is too risky	○ 1	○ 2	○ 3	○ 4	○ 5
In general, I believe a positively skewed distribution is riskier than a negatively skewed distribution	○ 1	○ 2	○ 3	○ 4	○ 5
In the first part of the experiment, I based my decisions more on the probability of gains and losses than on the pictures of the distributions	○ 1	○ 2	○ 3	○ 4	○ 5

Figure 9: Questionnaire

Appendix B - Skewness comparability and stochastic dominance

Experiments eliciting preferences over distributions differing in skewness cannot allow inference over the sign $U'''(w)$ for the reason indicated in Brockett and Kahane (1992) and Brockett and Garven (1998), so skewness alone cannot order distributions based on downside risk (Menezes et al., 1980). However, considering Menezes et al.'s definition of downside risk *"g(x) has more downside risk than f(x) if g(x) can be obtained from f(x) by a sequence of MVPTs"* and proposition 3b "*Let g(x) and f(x) be distributions with the same mean and variance. If f(x) dominates g(x) by TSD then g(x) can be obtained from f(x) by a sequence of MVPTs.*", then sufficient conditions for a distribution of having less downside risk than the other are (i) to have the same mean and variance, and (ii) to be third degree stochastic dominant. Moreover, respecting these conditions also implies being more right-skewed (Proposition 2a). These conditions are met by the distributions employed in our experiment. Hence, in our rounds *Binary-base*, *Multiple-base*, *Multiple-partial*, and

Gender		
	Males	88
	Females	91
	Other	1
Race		
	White	112
	Black/brown	47
	Latino	11
	Other	10
Education level (completed)		
	High School	53
	Bachelor	69
	Master	45
	MBA/PhD	6
	Other	7
Education field		
	STEM	110
	Non-STEM	70
Profession		
	Student	107
	Part-time worker	14
	Full-time worker	49
	Other	4
Continent		
	Europe	110
	Africa	49
	Americas	14
	Asia	6
	Oceania	1
Age		
	Mean	24.86
	Median	24
	St. dev.	3.68

Table 4: Descriptive statistics about the subjects in the sample

Multiple-full a distribution with a higher skewness coefficient is (a) a downside risk decrease, and (b) more right-skewed. In the *Skew-risk*, and in the *Skew-risk-reference*, these statements are true for all sets of distributions with the same standard deviation.

Skewness comparability criteria

Following Chiu (2010), we say that two distributions F and G are skewness comparable in the sense of Van Zwet (1964), if the function $F^{-1}(G(x))$ is either convex or concave. If $F^{-1}(G(x))$ is convex, then F is more positively skewed than G. Oja (1981) provided a weaker version of skewness comparability: two distributions F and G are skewness comparable in the sense of Oja if $F(\sigma_F x + \mu_F)$ and

$G(\sigma_G x + \mu_G)$ cross exactly twice. F is more positively skewed than G if G crosses F exactly twice, first from above. Chiu (2010, definition 5) also provides a more general form of skewness comparability, which relates to Menezes et al.'s definition of downside risk: *"distributions F and G are (generalized) skewness comparable if* $[F(\sigma_F x+\mu_F) \to G(\sigma_G x+\mu_G)]$ *is a downside risk increase or downside risk decrease or* $F(\sigma_F x + \mu_F) = G(\sigma_G x + \mu_G)$*"*, with the expression in the brackets interpreted as passing from distribution F to distribution G.

Chiu's (2010) second lemma posits that Van Zwet's comparability implies Oja's comparability, and Oja's comparability implies generalized skewness comparability. We claimed that, in our experiment, choosing a distribution with the same mean, same variance, and lower (higher) skewness, was a downside risk increase (decrease). To support this claim, proof of either of the three definitions of comparability is enough. Here we show Oja's criterion is met since it is visually easier: figure 10 reports three pairwise comparisons of the distribution functions of the "standardized distributions" of the three distributions employed in the *Binary-base* round. In all three plots, the distribution functions cross each other twice, and the distribution with the lower skewness coefficient (G) crosses that with a higher coefficient (F) first from above. We do not report here every possible pairwise comparison, but the graphs would be equivalent for any pair of distributions differing in skewness (and having the same mean and variance).

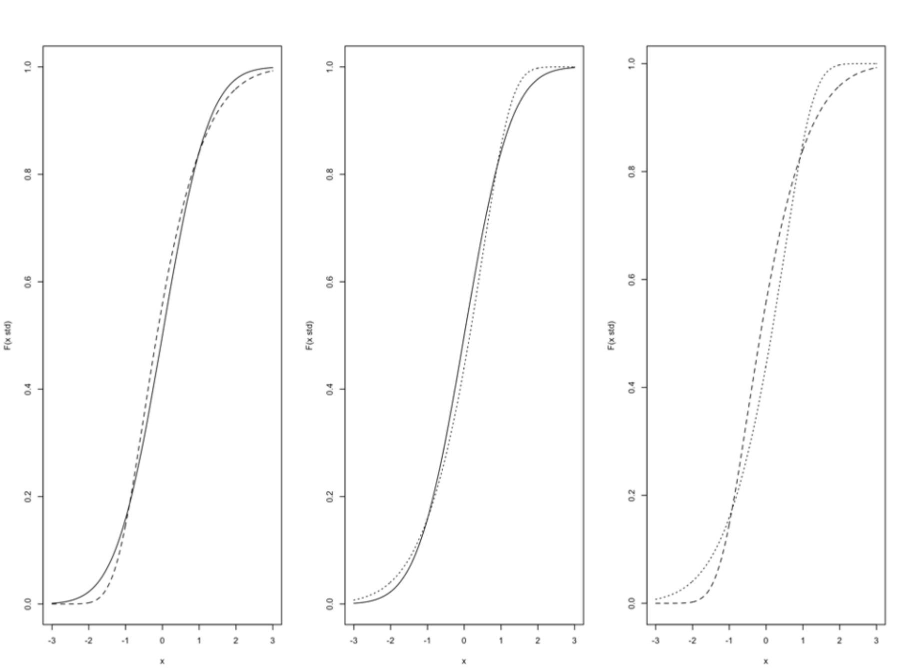

Figure 10: Verification of Oja's skewness comparability criterion for distributions with skewness coefficient equal to -0.75, 0, and 0.75. The symmetric distribution is reported with the continuous line, the positively skewed distribution with the dashed line, and the negatively skewed distribution with the dotted line.

Stochastic dominance

In rounds *Binary-base*, *Multiple-base*, *Multiple-partial*, and *Multiple-full* subjects chose a distribution from a given set of distributions, none of which was first-order or second-order stochastically dominated. However, the distributions with a higher skewness coefficient third order stochastically dominated (TSD) those with a lower skewness coefficient. If the decision maker utility function is such that $U^{'}(w) > 0$, $U^{''}(w) < 0$ and $U^{'''}(w) > 0$, then the dominant alternative has a higher expected utility (Levy, 1992). In rounds, 2 and 3, the distributions were the same as round 1, but with an adjustment of the parameter ξ, which increased or decreased the expected return of one alternative by up to 1%. The introduction of the bonus and penalty may have altered the situation of stochastic dominance with respect to the first round depending on the treatment. In rounds 7 and 8, subjects chose among distributions differing in skewness and standard deviation. In these two rounds, all distributions such that standard deviation was not minimized were second-order stochastically dominated by the other distributions with the same expected return, the same skewness but lower variance. If the decision maker utility function is such that $U^{'}(w) > 0$, $U^{''}(w) < 0$, then the dominant alternative has a higher expected utility (Levy, 1992).

Given two random variables A and B, with distribution functions $F_A(x)$ and $F_B(x)$, then

- "A" first-order stochastically dominates "B" if

  $F_B(x) - F_A(x) \geq 0 \, \forall x$ with strict inequality for some x.

- "A" second-order stochastically dominates "B" if

  $\int_{-\infty}^{x} [F_B(t) - F_A(t)]dt \ \geq 0 \, \forall x$ with strict inequality for some x.

- "A" third-order stochastically dominates "B" if

  $\int_{-\infty}^{x} \left[\int_{-\infty}^{\tau} [F_B(t) - F_A(t)]dt \right] d\tau \ \geq 0 \, \forall x$ with strict inequality for some x.

Figures 11, 12, and 13 report the graphs of the three functions above for the distributions employed in the *Binary-base* of the experiment. For all three plots, "A" is the alternative with the larger skewness coefficient, and "B" is the alternative with the smaller skewness coefficient.

In all three figures, the first plot, which represents $F_B(x) - F_A(x)$, lay both above and below the horizontal axis (dashed line), indicating that neither distribution first-order stochastically dominates the other. The same applies to the second plot, which represents $\int_{-\infty}^{x} [F_B(t) - F_A(t)]dt$, indicating that neither distribution second-order stochastically dominates the other. The third plot, which represents $\int_{-\infty}^{x} \left[\int_{-\infty}^{\tau} [F_B(t) - F_A(t)]dt \right] d\tau$, lay above the horizontal axis in all three figures, meaning that A third-order stochastically dominates B for all three pairwise comparisons. In our experimental framework, for all pairwise comparisons between distributions with the same mean and variance, the distribution with the larger skewness coefficient TSD and the other one with a lower skewness coefficient.

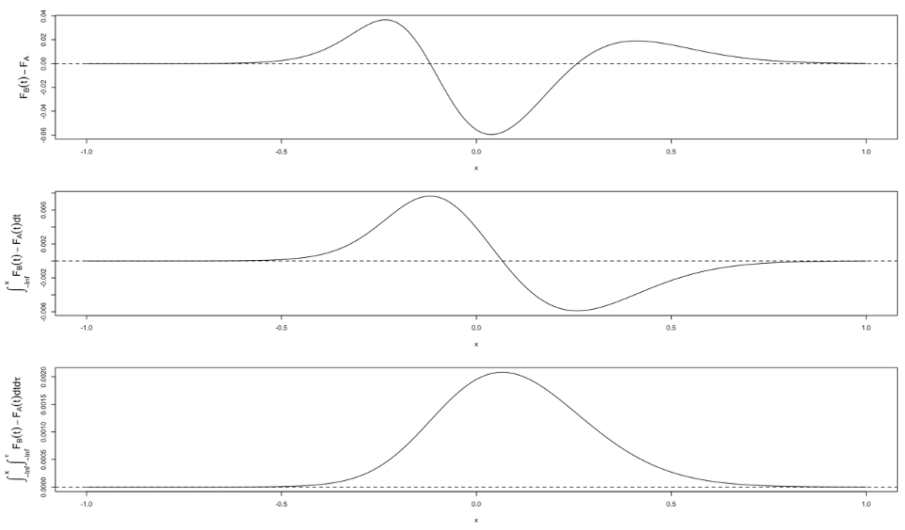

Figure 11: Stochastic dominance tests for the positively skewed distribution (A) and the symmetric distribution (B)

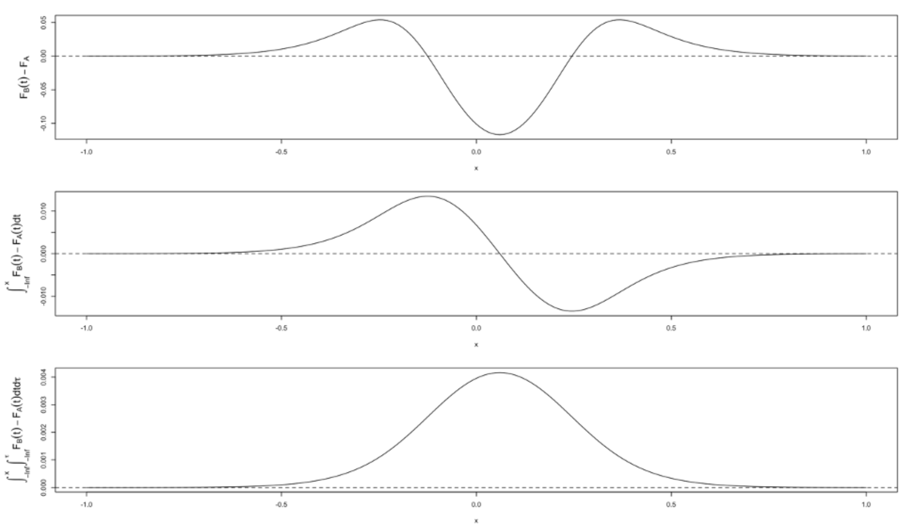

Figure 12: Stochastic dominance tests for the positively skewed distribution (A) and the negatively skewed distribution (B)

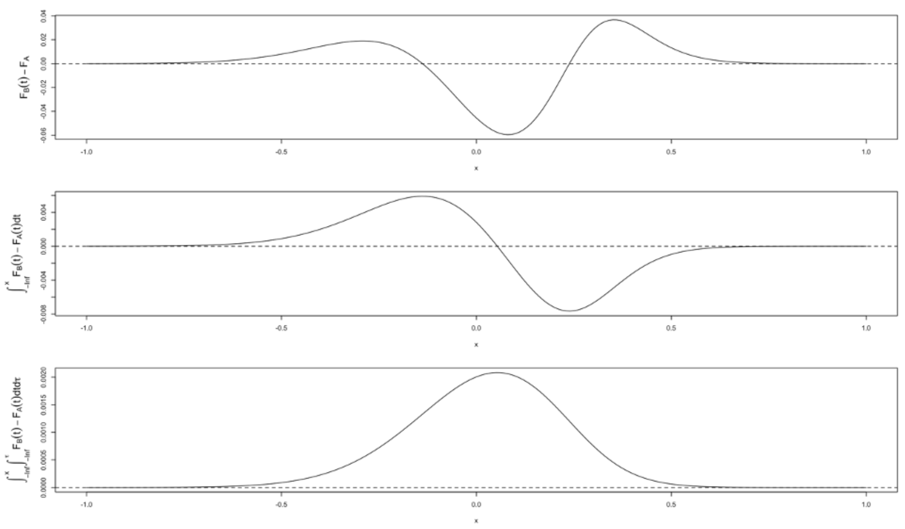

Figure 13: Stochastic dominance tests for the symmetric distribution (A) and the negatively skewed distribution (B)

Appendix C - Additional analyses

Skewness and risk-taking

We have argued that the main reason for subjects assigned to the negative treatment to take more risk than those assigned to the positive treatment is to improve their upside. However, they should in theory first maximize skewness to improve their upside, and only later increase variance. Therefore, we should expect a positive and significant correlation between skewness and variance also in the *Skew-risk* round, at least in the negative treatment. While a simple statistical test on choices suggests the correlation is positive but insignificant, a deeper analysis shows a situation more in line with our expectations: for each of the two treatments, we split the subjects into two groups based on the selected skewness level: those subjects choosing a distribution with a skewness coefficient lower than 0.5 (in absolute value) were in the low skewness group, and the others in the high skewness group. We find that both the subjects in the low negative skewness, and in the high positive skewness took more risk than the subjects assigned to the same treatment, but in the opposite "skewness group" (Figure 14). In the *Skew-risk* round, this difference is significant only in the negative treatment, while in the *Skew-risk-reference* round, it is significant for both treatments and at a 1% level (Figure 15).

Skewness and prudence

In Table 5 we report the proportions of subjects exhibiting prudent behavior in the *Binary-base* round (for each treatment), and skewness-seeking behavior in the *Multiple-base* round. In theory, this round could be also used to test for prudence: a prudent decision-maker should have selected the distribution with the highest skewness coefficient. About 17% of the subjects selected this distribution, signifi-

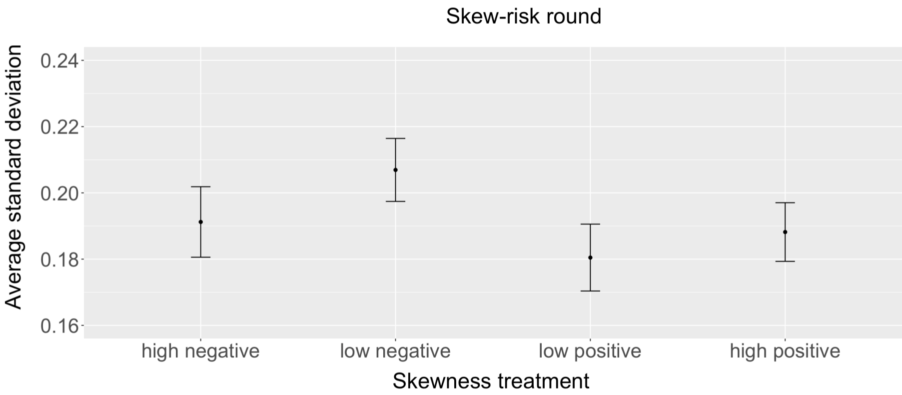

Figure 14: Choices in the *Skew-risk* round with division based on treatment and skewness level chosen. The subjects choosing distributions with a longer right tail took more risk. This is especially true for the negative treatment.

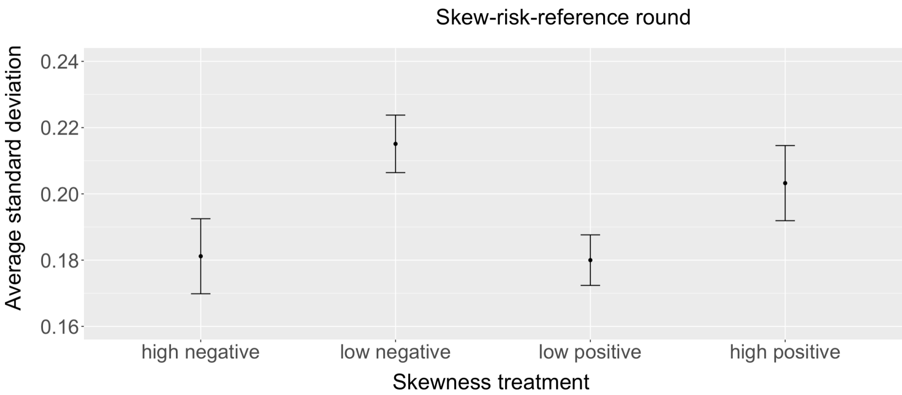

Figure 15: Choices in the *Skew-risk-reference* round with division based on treatment and skewness level chosen. The subjects choosing distributions with a longer right tail took more risk. This is true for both treatments, and the effect is more pronounced than in the *Skew-risk* round.

cantly more than what would be implied by random choice ($p < 0.001^{***}$). While the round was not designed with the purpose of testing for prudence, this result confirms the finding that the direction of skewness matters more than its absolute value (Brünner et al., 2011; Ebert, 2015). Indeed both skewness-seekers and skewness-avoiders did not cluster on the most positively skewed or the most negatively skewed distribution, but they selected several distributions, both with high and low absolute skewness.

Robustness checks on determinants of skewness preferences

In Table 3 we modeled the skewness levels chosen by the subjects in the three *Multiple* rounds. We identified the risk perception and the speculative channels as the main drives of skewness choices. In some specifications, we identified the

Round	Proportion of skewness-seeking/ prudent choices	p-value
Binary-base (pooled)	44.44%	0.16
Treatment pos and symm	36.84%	0.06
Treatment neg and symm	47.83%	0.81
Treatment pos and neg	48.15%	0.89
Multiple-base	66.11%***	<0.001

Table 5: Results of the Chi-squared test on proportions, Ho: proportion of skewness-seekers is equal to 50%

speculative channel through a division of the sample using the call/put preference. Furthermore, we divided the downside-focused subjects based on the maximum loss threshold. While the first criterion about the call/put preferences did not require any additional assumption, the second criterion required the specification of a loss threshold. The choice of 5% and 10%, although motivated by previous literature and by our data, was somewhat arbitrary. If we changed the threshold for classification, choosing a larger value (up to 15%), the coefficient of risk perception would not change in magnitude and still be highly significant, and the coefficient of the 'Call group' would not change in magnitude and be still significant, either at 1% or 5% level, depending on the set threshold. Hence, the two channels we identified are robust to the threshold chosen for the maximum loss.

However, as we increase the threshold from 5% to 15%, the coefficient of the 'Put-risk' group would reduce (still remaining positive in all ten alternative specifications) and lose its significance at a 5% level. This is due to the fact that as we increase the threshold for losses, the group 'Put-safe' starts to include subjects who have relatively higher thresholds for losses.

Consistency in choices

We show that subjects were consistent in their choices with two approaches. First, we estimated SUR using the choices of all rounds except from the two *Binary-adjustment* (because the expected return was different for the alternative distributions). Here we report the correlation coefficient of the residuals of the 6 regressions and the p-values of the statistical tests. All residuals are positively correlated, and this correlation is statistically significant in most cases.

	Bin-base	Mult-base	Mult-part	Mult-full	SR	SRR
Bin-base	1	0.24**	0.21**	0.12	0.13	0.07
Mult-base		1	0.40***	0.24**	0.24**	0.28***
Mult-part			1	0.26***	0.27***	0.27***
Mult-full				1	0.51***	0.30***
SR					1	0.33***
SRR						1

Table 6: Correlation across residual of SUR

Secondly, we use a simulation approach. We compare the standard deviation of the skewness level of the choices in the six rounds indicated above with the standard deviation when these choices were random. If subjects were consistent across the rounds, the standard deviation of the chosen skewness levels should be lower than if the decisions were random. The simulation shows that the median actual standard deviation is significantly lower than the simulated randomized standard deviation (p-value of Wilcoxon rank sum test with continuity correction is $< 0.001^{***}$), confirming consistency in choices. Figure 16 shows that real choices have more density located in the low standard deviation area than simulated choices.

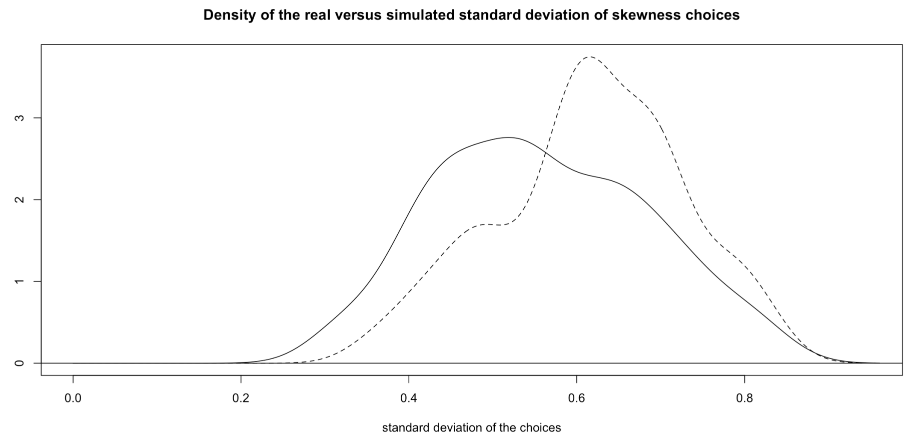

Figure 16: Estimated density of the standard deviation of the skewness level of rounds *Binary-base*, *Multiple-base*, *Multiple-partial*, *Multiple-full*, *Skew-risk*, *Skew-risk-reference* (solid line), and simulation of the same density if choices had been random (dashed line).