

Ensthaler, Ludwig; Nottmeyer, Olga; Weizsäcker, Georg

Working Paper

Hidden skewness

DIW Discussion Papers, No. 1238

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: Ensthaler, Ludwig; Nottmeyer, Olga; Weizsäcker, Georg (2012) : Hidden skewness, DIW Discussion Papers, No. 1238, Deutsches Institut für Wirtschaftsforschung (DIW), Berlin

This Version is available at:

<https://hdl.handle.net/10419/66860>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

1238

Discussion Papers

Deutsches Institut für Wirtschaftsforschung

2012

Hidden Skewness

Ludwig Enthaler, Olga Nottmeyer and Georg Weizsäcker

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2012

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN print edition 1433-0210
ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Hidden skewness ^{*}

An earlier version of this paper was circulated as DIW DP 1043

Ludwig Ensthaler Olga Nottmeyer Georg Weizsäcker

Abstract

Multiplicative growth processes that are subject to random shocks often have a skewed distribution of outcomes. A simple laboratory experiment shows that participants either strongly underestimate skewness or ignore it completely. The participants' choices reveal bounds on their subjective medians of a financial asset's price that is subject to stochastic growth. The observed bias in expectations is irrespective to risk preferences and fairly robust to feedback. It is consistent with a behavioral model in which geometric growth is confused with linear growth. The bias is a possible explanation of investors' misunderstandings of real-world financial products like leveraged ETFs.

JEL-Classification: C91, D03

Keywords: skewness, belief biases, binomial tree

^{*}Ensthaler: University College London, l.ensthaler@ucl.ac.uk; Nottmeyer: IZA, nottmeyer@iza.org; Weizsäcker: University College London and DIW Berlin, gweizsaecker@diw.de. We are grateful for helpful comments by Erik Eyster, Dorothea Kübler, Erik Mohlin, Peter Mörters, Tobias Schmidt, Adam Szeidl, Heinrich Weizsäcker and several conference audiences, as well as for excellent research assistance by Mark Henninger. Weizsäcker thanks the ELSE Research center at University College London and the European Research Council (ERC Starting Grant 263412) for financial support.

1 Introduction

In the experiment presented in this paper, participants consider a security that has a seemingly simple price transition. They are told that the security, if bought, has to be held for exactly 12 months and is then to be sold.

You can buy the security at a price of 10,000 Euros. During each month, the security's price either increases by 70% or decreases by 60%. The two possible price changes in each month occur with equal probabilities ("fifty-fifty").

The instructions also explain that all random draws are independent. They are written for maximal clarity, with the important exception that they do not show the values of any compound price changes that accumulate over time. The participants may thus misperceive the random price process, given its compound nature.

The actual distribution of the security's selling price is, as the reader can verify, extremely skewed. A decrease by 60% cannot be undone by a single increase by 70% and the typical price path therefore tends downward. If the security was held infinitely long rather than 12 months, the price would converge to zero in probability. But already with a fixed maturity after 12 months, the *median* selling price is as low as 989 Euros. Skewness shows in the observation that the *mean* selling price after 12 months is much higher, at 17,959 Euros—the fact that 70% exceeds 60% implies that, in expectation, increases dominate decreases.

Our laboratory experiment tests whether the participants correctly locate the median. Through a sequence of simple choice problems we identify bounds on the median of each participant's subjectively expected distribution and find that it is typically far too high: 98% of the participants reveal that they have a subjective median above 2,000 Euros, and 84% above 9,000 Euros. We conclude that the participants have an incomplete understanding of the compound effects of multiplicative shocks in our example. The data is consistent with a behavioral model presented in Section 2, where

the decision-maker mistakenly views any multiplicative changes as additive changes. A further result is that the effect is fairly robust to learning from feedback.¹

One possible economic consequence of the bias shows in the market of a relatively new class of investment products, so-called leveraged Exchange Traded Funds (ETFs). Leveraged ETFs pay investors a multiple of the daily percentage changes of the underlying indices that they track. For example, a triple ETF on the Dow Jones Industrial Average rises by 3% if the index itself increases 1% on a certain trading day. Such products can successfully serve short-term hedging objectives but if an investor holds the product for an extended period of time, she may be surprised by the high likelihood of making a loss that arises from multiplicative compounding. In particular, the increased skewness in returns means that leveraged ETFs can make a loss although the underlying index makes a profit. Accurate investor expectations may be especially doubtful as significant proportions of leveraged ETFs are held by household investors.² Not long after their introduction leveraged ETFs came under scrutiny. The Financial Services Regulatory Authority published a caution notice and severe anecdotal evidence quoted in newspaper articles and investor platforms suggests that many private investors misunderstand the risks of leveraged ETFs.³ While a number of unrelated factors such as a complex fee structure could contribute to the frequent complaints, we conjecture that there is a more fundamental bias that drives the misconception of these products, namely people's failure to correctly compound multiplicative growth processes.

The cognitive difficulties of multiplicative compounding has also been documented in various other settings. The literature on exponential growth bias (Eisenstein and Hoch 2005, Stango and Zinman, 2009, among others), finds that people underestimate interest accumulation and related growth

¹In the fifth repetition of our experiment, with feedback about the realized selling prices (detailed in Section 2), 86% of subjective medians are still above 2,000 Euros and 55% above 9,000 Euros.

²See the assessment of Brendan Conway of Barron's at <http://blogs.barrons.com/focusonfunds/2012/03/21/who-uses-leveraged-and-inverse-etfs-anyway/> (accessed 16/08/2012).

³See e.g. *Wall Street Journal* (2012).

processes. For instance, Chen and Rao (2007) document how retailers can strategically use multiple price discounts, as many consumers may interpret a double dip of 20% followed by another 25% as a total reduction of 45%, instead of the correct 40%. But the exponential growth bias is usually studied in the context of deterministic growth, whereas we study random returns. Here, important previous evidence is in Benartzi and Thaler (1999) who show that people choose different (hypothetical) retirement saving plans depending on whether the framing of the decision problem requires them to calculate the distributions arising from compound random changes. A context even closer to our study are compensation packages with employee stock options, where an employee receives the right to purchase a certain number of shares at a given price at a certain date in the future. If the employee misperceives the probability distribution of the share price at maturity, she may accept a suboptimal package. In particular, this may happen if she ignores the skewness and perceives a too-fat right tail of the share price distribution.⁴

Our experiment matches the simplest work horse model in the option pricing literature, the so-called binomial tree model of Cox, Ross and Rubinstein (1979). There, the price of the underlying asset increases or decreases by known factors and with known probabilities, just like in our experiment, and arbitrage-free option prices are derived. Adding to the scant experimental evidence on misperceptions of financial options (e.g. Gneezy, 1996, Abbink and Rockenbach, 2006), our results show a particular bias in the perception of this basic model: at least with our set of parameters, the typical price path appears in far too positive a light. Independent from our work, Stutzer and Jung Grant (2010) analyse investment choices with essentially the same asset. Their results are consistent with ours in that they find an inflated investment rate if participants have to calculate compound changes

⁴Stock options are particularly popular with start up companies who have not got the financial resources to pay competitive salaries (yet). Importantly, start up companies are also among the employers with the highest proportional per-period volatility in market value, and thus with high skewness of long-run share price.

themselves.⁵⁶

The next section contains the behavioral model and details on the experiment, while the appendix contains the full instructions. Section 3 shows the results and Section 4 concludes.

2 Experimental design

The model: Following the described evidence of linearized perceptions, we model a biased decision maker who obtains a subjective distribution as follows. First, the decision maker correctly computes the distribution of changes in the relevant variables occurring in $t = 1$. Second, she is biased in the sense that she uses linear extrapolation to extend the distribution of absolute changes to T periods. Formally, let Y_0 denote the initial price of an asset and let μ_t be the random variable describing the relative price changes occurring in t , e.g. $Y_1 = Y_0 \mu_1$. Let η_t denote the random variable describing the absolute price movements in t . As in the basic model of Cox, Ross and Rubinstein (1979) we consider a price movement where μ_t has a distribution that is constant across t . The hypothesized bias is that the decision maker instead views the distribution of η_t as constant across t . That is, where an unbiased decision maker perceives the true distribution as $Y_T = Y_0 \prod_{t=1}^T \mu_t$ with $\{\mu_t\}$ i.i.d., the biased decision maker perceives the final price as $\tilde{Y}_T = Y_0 + \sum_{t=1}^T \eta_t$, with $\{\eta_t\}$ i.i.d. and its distribution equal to that of η_1 .

⁵But different from our experiment, theirs uses a hypothetical-choice context without actual monetary incentives. Another notable difference is that our experiment is designed to detect the bias without auxiliary assumptions on risk preferences.

⁶Our experiment is also reminiscent of additive random processes that have been studied experimentally, see e.g. Redelmeier and Tversky (1992) Benartzi and Thaler (1999) and Klos, Weber and Weber (2005), all of whom follow up on Samuleson's (1963) hypothetical offer of a sum of gambles to his colleague. While a formal connection to our security appears immediately by taking the logarithm, a key difference is that taking the logarithm of our security leads to a sum of less-than-fair gambles, not more-than-fair gambles like those studied in previous experiments. This feature of our security is equivalent to the property that a 60% decrease weighs proportionally stronger than a 70% increase, creating the extreme skewness that we examine.

It is straightforward to calculate properties of the perceived distribution, e.g. its mean $E[\tilde{Y}_T] = Y_0 + TE[\eta_1]$. If η_1 is symmetrically distributed, as in our experiment, \tilde{Y}_T is symmetrically distributed, too. Thus, mean and median coincide.⁷

Choice problems: The experiment is designed to elicit the participants' expectations, irrespective of their risk preferences. The monetary incentives in each choice problem therefore involve only two possible payments—"receive a bonus" versus not—making it optimal for any decision-maker with monotonic preferences to maximize the subjectively perceived probability of receiving the bonus.

The choice problems are framed in a financial investment context: two risky securities are offered and the selling price of the chosen security determines whether the bonus is paid.⁸ Security *A* is the security described in the introduction. A participant who chooses this security receives the bonus if the selling price at maturity exceeds a given threshold t_A . The alternative choice is Security *B*, which yields the bonus with probability one half. One can immediately see that it is subjectively optimal for a participant to choose Security *A* if and only if she believes that Security *A* yields the bonus with probability more than one half. A choice for Security *A* thus reveals that the median of her subjective probability distribution of Security *A*'s selling price is above t_A .

For a balanced description of the two choice options, Security *B* is phrased analogously to Security *A*, with the difference that only a single price change of +70% or -60% (equiprobably) occurs during the 12 months. A participant who chooses Security *B* receives the bonus if the selling price of *B* exceeds a separate threshold t_B . This threshold is fixed at the initial price of 10,000

⁷The model is consistent with the biased perceptions of deterministic multiplicative processes in the literature, such as compound interest accumulation. Let e.g. $Y_T = Y_0(1+r)^T$, where r is the annual rate of interest. The linearization bias as described above has the decision maker perceive a constant additive increase instead, $\tilde{Y}_T = Y_0 + T\tilde{r}$, where $\tilde{r} = Y_1 - Y_0$. For $r > 0$ and $T > 1$, we have $\tilde{Y}_T < Y_T$, consistent with the lab and field data that has been collected (e.g. Stango and Zinman, 2009).

⁸The descriptions begins with the wording: "You are a manager and have to make a decision between two risky investments."

Euros throughout the experiment whereas the threshold t_A varies between 10 different values (ranging from 100 to 250,000 Euros). Each experimental participant makes a choice between A and B for each of the 10 possible values of t_A , allowing us to infer bounds on her subjective median of the selling price of Security A . Table 1 shows the 10 choice problems as seen by the participants.

Table 1: The 10 binary choices

	Threshold for Security A	Threshold for Security B	Your decision (A or B)
Task 1	100	10,000	---
Task 2	500	10,000	---
Task 3	2,000	10,000	---
Task 4	6,000	10,000	---
Task 5	9,000	10,000	---
Task 6	12,000	10,000	---
Task 7	20,000	10,000	---
Task 8	35,000	10,000	---
Task 9	90,000	10,000	---
Task 10	250,000	10,000	---

Treatment conditions: Participants are randomly assigned to one of two conditions that differ in the extent to which the experimental instructions explain the implied distributions. The CONTROL condition presents the basic explanation. To introduce Security A , the instructions use the above formulation “You can buy...”. This is followed by a statement about the independence of random draws and by the paraphrase that after month 1, the security’s price is either at 17,000 Euros or at 4,000 Euros. The instructions then repeat the random price transition, but without calculating compound effects explicitly: “At the end of month 2, the price is either 70% higher or 60% lower than at the end of month 1. At the end of month 3, the price is either 70% higher or 60% lower than at the end of month 2. And so on, ...” Security B is introduced with identical wording to that of Security A (where applicable). Next, the thresholds t_A and t_B are explained and two examples are given. Finally, participants face an understanding test of four questions which they have to answer correctly before they may proceed. The examples

and understanding test are carefully chosen to not suggest any responses to the participants.

A possible concern is that data patterns in the CONTROL condition are not due to cognitive limitations but instead are driven by the choice format, the context frame or other cues. In particular, the set of 10 threshold values can conceivably influence the responses.⁹ We address this concern by including the TREATMENT condition where we provide the participants with an additional explanation, leaving the remainder of the instructions unchanged. The additional text (about one written page) gives an explicit calculation of the distribution of compound price changes after two periods. It also points out the asymmetry in the selling price distribution and lists the implicit probabilities of receiving the bonus from choosing Security *A* for each value of t_A . Any difference in responses under the two conditions must stem from differences in the understanding of these implied truths.

Feedback and repetitions: After the participants make their 10 choices, they receive individual feedback in the form of a sample pair of selling prices of Securities *A* and *B*. This concludes the first round of the experiment. The experiment is then repeated for four additional rounds of the same nature, each including 10 choices and individual feedback. The feedback procedure and the choice format are identical for both treatment conditions.¹⁰

Procedures and payments: All 128 participants (68 in CONTROL and 60 in TREATMENT) are students at Technical University Berlin. Six sessions, three in each treatment condition, are conducted in a paper-and-pencil format. The protocol is fixed across all sessions. The instructions are read aloud to the participants, up to the beginning of the understanding test. Participants receive a participation fee of 5 Euros and a possible bonus of 5 Euros per round. That is, participants can earn up to five bonuses of 5

⁹We deliberately fixed the 10 values of t_A so that half of them exceed Security *A*'s starting price of 10,000 Euros, in order to not suggest a direction of price change. However, this property may conceivably induce a midpoint effect, leading the participants to switch from *A* to *B* towards the middle of the list.

¹⁰Each additional round comes with the chance to earn a new bonus (see the next paragraph in the main text), but this does not affect the simple optimality conditions for choice. Independent of other choices it remains optimal to choose *A* iff the subjective median is above t_A .

Euros each, one per round of the experiment. After completing all choices, each participant receives five random draws of integers between 1 and 10 to determine which of the 10 choice problems in each round is payoff relevant for her. She receives the bonus for a given round if the selling price of the chosen security in the payoff-relevant problem exceeds its threshold.

3 Results

The data analysis is simplified by the observation that a participant with any subjective belief about selling prices maximizes her preference by choosing Security A for low values of t_A and switching to B for all values higher than her subjective median, i.e. she switches between the securities no more than once. We observe such unique switching points in the large majority of responses (93%) and restrict attention to these data.¹¹

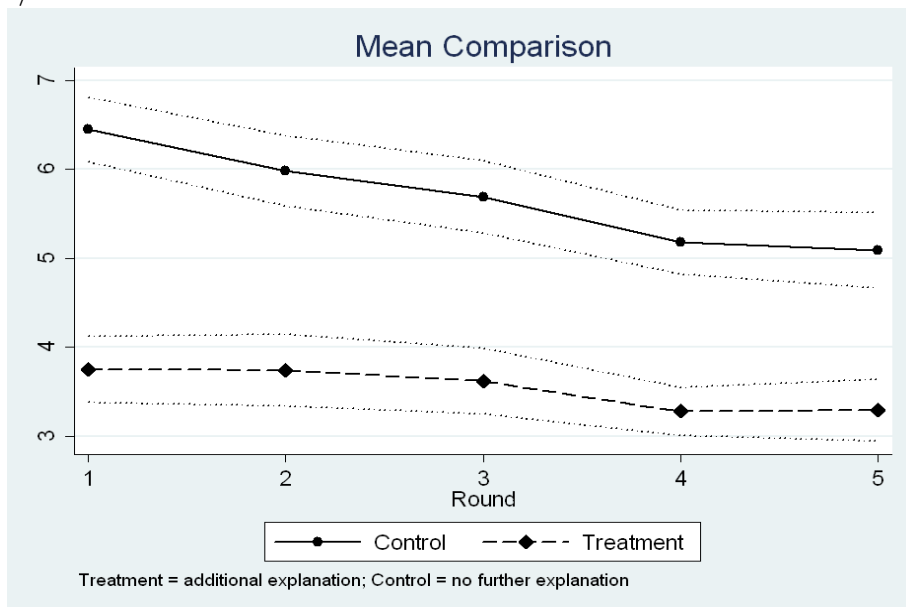
The benchmark “rational” prediction is for all participants to choose A in the first two tasks and then switch to B . This is optimal as the true median of A ’s selling price is between the threshold values of Task 2 and Task 3. The above behavioral model predicts a higher switching point: with perceived selling price \tilde{Y}_T , i.e. a perceived median of 16,000, the decision maker chooses A in the first six tasks and then switch to B .

Figure 1 shows the mean switching point for each round of the experiment, separately for CONTROL (solid line) and TREATMENT (dashed line). More precisely, it shows the mean of task numbers at which participants start choosing Security B .¹² The dotted lines show the 95% confidence intervals, taken pointwise around the means at each round of the experiment. As shown in the figure, the mean switching point in CONTROL is 6.5 in the first round of the experiment, close to the switching point of 7 as predicted by the behavioral model. It decreases to 5.1 in the fifth round of the experiment. In TREATMENT, the mean switching point is 3.8 in round 1 and decreases to

¹¹If a participant has multiple switching points in one round, her answers in the remaining rounds are still considered. None of our conclusions would change if we dropped all responses by subjects who switch strictly more than once in at least one round (12% of participants), or if we included all data and considered each of the 10 tasks separately.

¹²We assign the value 11 if a participant always chooses A .


Figure 1: Means of switching points, separated by round and CONTROL/TREATMENT.



3.3 in round 5, close to the rational prediction. Figure 2 shows histograms of the distributions of switching points, again separately for each round of the experiment and for the two conditions CONTROL and TREATMENT. Both figures show strong differences between the two conditions, and parametric t-tests as well as non-parametric Wilcoxon rank-sum tests confirm that all round-by-round comparisons between the two conditions are statistically significant at $p < 0.001$. In particular, the treatment effects are still highly significant in the last round of the experiment.

Table 3 reports the distributions of switching points that underlie Figure 2 and lists the implied ranges for the medians of the participants' subjective distributions of Security A 's selling price. Not a single CONTROL participant in round 1 reveals a subjective median between 500 and 2,000 Euros (i.e. rational switching at Task 3). Instead, 98% of CONTROL participants reveal that their subjective medians are above 2,000 Euros in round 1, and still 86% in round 5. The modal choice in round 1 is a switching point of 6 indicating a subjective median between 9,000 and 12,000 Euros. Roughly 20% of CONTROL participants choose a switching point at Task 7. This

Figure 2: Distribution of switching points, separated by round and CONTROL/TREATMENT.



reveals subjective medians between 12,000 and 20,000 Euros, consistent with the point prediction of our behavioral model. Further 23% of participants choose an even higher switching point. Under the TREATMENT condition, 70% of responses are at the optimal switching point of 3 already in round 1. Altogether, the data show a consistent pattern that the performance is poor under the CONTROL condition—close to the behavioral model’s prediction—and much better in TREATMENT.

We run random effects regressions to obtain a better description of responses over time, exploiting the panel structure of the data. This allows accounting for individual heterogeneity as well as describing the reaction of participants to their individually different feedback information.¹³ The dependent variable is a participant’s observed switching point in a given round,¹⁴ and the explanatory variables are *Treatment* (1 if in condition

¹³The appropriateness of random effect regressions is confirmed by applying a Hausman-test that does not reject the statistical independence between unobserved factors and the explanatory variables used. A comparison to pooled OLS regressions shows no substantial differences across comparable coefficients.

¹⁴Translating the task number into the corresponding subjective median would not

TREATMENT, 0 if in CONTROL), *Round* and *Feedback*. The latter is a dummy variable that is 1 if the participant’s sample feedback in the previous round has the property that Security *B*’s selling price exceeds that of Security *A*. In this case, participants get the ‘correct’ feedback that returns to investment in Security *A* are likely to be small.

Table 2: Results from Random Effects Regressions

Dep. Var: Switch	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Treatment</i>	-2.129*** (0.19)	-2.123*** (0.19)	-2.791*** (0.33)	-2.136*** (0.19)	-2.095*** (0.20)	-2.135*** (0.19)	-2.755*** (0.33)
<i>Round</i>		-0.237*** (0.04)	-0.356*** (0.06)			-0.265*** (0.04)	-0.385*** (0.06)
<i>Treatment</i> × <i>Round</i>			0.223** (0.08)				0.229** (0.08)
<i>Feedback</i>				-0.248* (0.12)	-0.355* (0.14)	-0.462*** (0.12)	-0.610*** (0.15)
<i>Treatment</i> × <i>Feedback</i>					-0.187 (0.24)		-0.299 (0.23)
<i>Constant</i>	5.683*** (0.14)	6.393*** (0.19)	6.748*** (0.23)	5.879*** (0.17)	5.963*** (0.18)	6.840*** (0.22)	7.316*** (0.28)
<i>N</i>	596	596	596	596	596	596	596

The estimation of model (1) in Table 2 repeats the main result that the additional explanation in TREATMENT has a significant effect. Comparisons with the richer models show that the coefficient is fairly robust to changes in the specification. The coefficient of *Round* is negative and significant (model (2)) indicating that participants adjust their decision over time. Moreover, participants in CONTROL make greater progress across rounds, as shown in model (3). There, a test for sums of coefficients shows that participants in condition TREATMENT do not significantly change their response over time.

Regarding the participant’s reaction to feedback, the coefficient of the *Feedback* dummy variable has the expected negative sign, i.e. participants switch from A to B at a lower threshold if their feedback shows a relatively low selling price for Security *A*. The effect is less significant, however, if *Round* is not included (models (4) and (5)).¹⁵ Models (6) and (7) include

change the main conclusions. But the subjective medians have some extreme outliers, complicating the analysis.

¹⁵In an alternative specification of the *Feedback* dummy we assign the value 1 if the participant’s selling price of Security *A* lies above their subjective median in the previous round. (To generate this variable, we linearly interpolate each participant’s subjective median to be the arithmetic average of the revealed bounds.) However, the corresponding

both *Round* and *Feedback*. The results remain essentially the same, except that the coefficient on *Feedback* is now significant at lower levels. Overall, the regression analysis confirms that participants in CONTROL have a poor understanding of the median selling price of Security A, whereas in TREATMENT their responses are significantly closer to the optimal response.

coefficient estimates are mostly insignificant.

Table 3: Subjective Medians

Switching point from A to B	Range of subjective Median for A	Share of people switching from A to B										
		Round 1		Round 2		Round 3		Round 4		Round 5		
		Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	
1	[0 - 100]	1.79	0.00	0.00	3.23	0.00	0.00	4.76	1.82	4.62	0.00	9.23
2	[100 - 500]	0.00	0.00	0.00	1.61	0.00	0.00	0.00	0.00	4.62	3.57	4.62
3	[500 - 2,000]	0.00	70.31	5.45	61.29	7.27	66.67	67.69	7.27	67.69	10.71	66.15
4	[2,000 - 6,000]	3.57	10.94	14.55	14.52	12.73	9.52	13.85	20.00	13.85	30.36	4.62
5	[6,000 - 9,000]	10.71	6.25	9.09	4.84	25.45	6.35	4.62	30.91	4.62	14.29	6.15
6	[9,000 - 12,000]	41.07	6.25	38.18	6.45	30.91	6.35	0.00	23.64	0.00	19.64	3.08
7	[12,000 - 20,000]	19.64	3.13	18.18	6.45	10.91	3.17	4.62	12.73	4.62	14.29	4.62
8	[20,000 - 35,000]	17.86	1.56	9.09	0.00	10.91	1.59	0.00	3.64	0.00	5.36	1.54
9	[35,000 - 90,000]	5.36	0.00	5.45	0.00	0.00	1.59	0.00	0.00	0.00	1.79	0.00
10	[90,000 - 250,000]	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
11	[250,000 - ∞]	0.00	1.56	0.00	1.61	1.82	0.00	0.00	0.00	0.00	0.00	0.00

4 Conclusion

Questions about compound interest are, by now, standard procedure in surveys about financial literacy—see e.g. the relevant module in the Health and Retirement Survey documented in Lusardi and Mitchell (2011). The typical evidence is that calculations of multiplicative growth effects show a strong downward bias, often to the extent that all compounding is ignored. The bias seems robust and economically important. However, very few studies include an incentivised experiment to corroborate the evidence (a notable exception from psychology is Christandl and Fetchenhauer, 2009) and no study has, to the best of our knowledge, included an alternative, behavioral model that formally captures the hypothesized false linearization of multiplicative processes. Our discussion in Section 2 of this paper aims to fill this gap. Such a model may be useful in that it can easily generate behavioral predictions for more general contexts, like in our experiment: if the growth process is not only multiplicative but also subject to randomness, a falsely perceived linearization of it may lead the decision-maker to miss out entirely on the skewness of the relevant distribution. Our experimental data demonstrates that the corresponding misperception is confirmed empirically. More generally (leaving the model aside), the data demonstrate that the misperception of skewness arising from random multiplicative growth can be sizable and that almost all experimental participants suffer from it.

References

- Abbink, K., and B. Rockenbach, 2006, Option pricing by students and professional traders: a behavioural investigation, *Managerial and Decision Economics* 27 (6), 497-510.
- Benartzi, S., and R.H. Thaler, 1999, Risk aversion or myopia? Choices in repeated gambles and retirement investments, *Management Science* 45, 364-381.
- Chen, H., and A.R. Rao, 2007, When two plus two is not equal to four: Errors in processing multiple percentage changes, *Journal of Consumer Research* 34, 327-340.

- Christandl, F., and D. Fetchenhauer, 2009, How laypeople and experts misperceive the effect of economic growth, *Journal of Economic Psychology* 30, 381-92.
- Eisenstein, E.M., and S.J. Hoch, 2005, Intuitive compounding: Framing, temporal perspective, and expertise, working paper, Johnson Graduate School of Management, Cornell University.
- Cox, J.C., S.A. Ross, and M. Rubinstein, 1979, Option pricing: a simplified approach, *Journal of Financial Economics* 7, 229-263.
- Gneezy, U., 1996, Probability judgments in multi-stage problems: experimental evidence of systematic biases, *Acta Psychologica* 93, 59-68.
- Klos, A., Weber, E.U., and M. Weber, 2005, Investment decisions and time horizon: risk perception and risk behavior in repeated gambles, *Management Science* 51, 1777-1790.
- Lusardi, A., and O.S. Mitchell, 2011, Financial literacy and planning: Implications for retirement wellbeing, NBER Working Paper 17078.
- Redelmeier, D.A., and A. Tversky, 1992, On the framing of multiple prospects, *Psychological Science* 3(3), 191-93.
- Samuelson, P. A., 1963, Risk and uncertainty: a fallacy of large numbers, *Scientia* 98, 108-113.
- Stango, V., and J. Zinman, 2009, Exponential growth bias and household finance, *Journal of Finance* 64 (6), 2807-2849.
- Stutzer, M., and S. Jung Grant, 2010, Expected return or growth rate? Choices in repeated gambles that model investments, working paper, Leeds School of Business, University of Colorado.
- Wall Street Journal, 2012, Beware of 'Leveraged' ETFs, 11/05/2012.

Appendix (not intended for publication): Instructions

Welcome!

You are about to participate in an experiment in decision making. Universities and research agencies have provided the funds for this experiment.

In this experiment we will first ask you to read instructions that explain the decision scenarios you will be faced with. We will also ask you to answer questions that test your understanding of what you read. Finally, you will be asked to make decisions that will allow you to earn money. Your monetary earnings will be determined by your decisions and by chance. All that you earn is yours to keep and will be paid to you in private, in cash, after today's session.

Only for coming here and completing the experiment, you will also receive a fixed participation fee of EUR 5.00. The earnings that you make during the experiment will be added to this amount.

It is important to us that you remain silent and do not look at other people's work. If you have any questions or need assistance of any kind, please raise your hand, and an experimenter will come to you. If you talk, exclaim out loud, etc., you will be asked to leave and will forfeit your earnings. Thank you.

[page break]

Procedure and payment structure

You are asked to make a sequence of decisions. There are five rounds in this experiment. In each round, you have the opportunity to earn a bonus of EUR 5.00. In what follows, the term “bonus” will always refer to these EUR 5.00. All bonuses that you earn in any of the five rounds will be paid to you in cash after the experiment.

Each round consists of a list of ten tasks. One of the ten tasks will be chosen by a random draw made on the computer. This task will be paid out for real. That is, if you were successful in the task that the computer picked, you will earn the bonus of EUR 5.00. If you were unsuccessful in the task that the computer picked, you will not receive a bonus in this round.

The tasks are described on the next pages.

[page break]

INVESTMENT TASK

Setting

You are a manager and have to make a decision between two risky investments, either to buy security A or to buy security B. Either security, if bought, has to be held for 12 months. After the 12 months you sell the security. Depending on your investment success, you have the chance to earn a bonus.

Security A:

You can buy the security at a price of 10,000 Euros. During each month, the security's price either increases by 70% or decreases by 60%. The two possible price changes in each month occur with equal probabilities ("fifty-fifty"). The direction of price change (increase/decrease) is not influenced by the direction of price changes in previous months.

Thus, at the end of month 1, the price is either 70% higher or 60% lower than at the beginning of month 1. That is, the price is either 17,000 Euros or 4,000 Euros. At the end of month 2, the price is either 70% higher or 60% lower than at the end of month 1. At the end of month 3, the price is either 70% higher or 60% lower than at the end of month 2. And so on, until you sell the security at its price at the end of month 12.

Security B:

You can buy the security at a price of 10,000 Euros. During month 1, the price of security B moves identically to the price of security A. After the end of month 1, the price stays constant until the end of month 12.

Thus, at the end of month 1, the price is either 70% higher or 60% lower than at the beginning of month 1, with equal probability. That is, the price is either 17,000 Euros or 4,000 Euros. The price then stays constant until you sell the security at the end of month 12.

The following rule determines your payment: If the selling price of the security that you bought is higher than a certain threshold, you receive the bonus.

[page break]

Thresholds

The thresholds differ between security A and security B. Security B's threshold always equals its initial price of 10,000 Euros. Security A's threshold varies between 100 and 250,000 Euros.

For each of the possible thresholds of security A and security B that are presented in the table below, you will be asked to make a decision between A and B. These are the 10 tasks for one round of this experiment.

	Threshold for security A:	Threshold for security B:	Your decision (A or B):
Task 1	100	10,000	
Task 2	500	10,000	
Task 3	2,000	10,000	
Task 4	6,000	10,000	
Task 5	9,000	10,000	
Task 6	12,000	10,000	
Task 7	20,000	10,000	
Task 8	35,000	10,000	
Task 9	90,000	10,000	
Task 10	250,000	10,000	

Example 1

Consider *Task 1*, where the threshold for security A is 100 Euros, and the threshold for security B is 10,000 Euros.

Suppose that you decide to buy *security A*. If the selling price of security A is higher than 100 Euros, you receive the bonus. If the selling price is less than or equal to 100 Euros, you do not receive the bonus.

Now, suppose instead that you decide to buy *security B*. If the selling price of security B is higher than 10,000 Euros, you receive the bonus. If the selling price of security B is less than or equal to 10,000 Euros, you do not receive the bonus.

Example 2

Consider *Task 2*, where the threshold of security A is higher than in the previous example, at 500 Euros, and the threshold for security B is again 10,000 Euros.

First, suppose that you decide to buy *security A*. In this case, if security A's selling price is higher than 500 Euros, you receive the bonus. Otherwise, you do not receive the bonus.

If, instead, you decide to buy *security B*, you receive the bonus if the selling price of security B is higher than 10,000. Otherwise, you do not receive the bonus.

And so on, analogously for *Task 3*, *Task 4*, etc., until *Task 10*.

[page break]

[The following page is for participants in condition *TREATMENT* only.]

How likely does security A's selling price exceed its threshold?

As security A's selling price is determined by 12 price changes, there are 13 possible selling prices for security A altogether: the lowest price results if security A's price decreases in each of the 12 months; the second-to-lowest price results if 11 price changes are decreases and 1 is an increase; and so on.

An important property of security A is that if the price decreases *once* it requires *multiple* price increases to compensate for the decrease. A single price increase by 70% cannot make up for a single decrease by 60%.

For example, consider the price at the end of month 2. If the price change in month 1 is downward, i.e. a decrease from 10,000 Euros to 4,000 Euros, then an increase in month 2 would only yield a price of 6,800 Euros, well below the starting price of 10,000 Euros. Likewise, if the first price change is an increase from 10,000 Euros to 17,000 Euros but the second price change is a decrease, then the price at the end of month 2 would again be only 6,800 Euros (which is 40% of 17,000 Euros). For the price to exceed 10,000 Euros at the end of month 2, the price would therefore have to increase twice in a row – from 10,000 Euros to 17,000 Euros in month 1, and from 17,000 Euros to 28,900 Euros in month 2.

The example illustrates a general feature of security A: it has a small probability of ending up at an extremely high price, and a large probability of ending up at low prices.

The following table shows how many price increases are required for security A's selling price to exceed the threshold, in each of the 10 investment tasks. The table's final column shows exactly how likely the selling price exceeds the threshold.

	Threshold for security A:	Required # of increases, to exceed threshold	Probability of exceeding threshold
Task 1	100	5 or more	80.6 %
Task 2	500	6 or more	61.3 %
Task 3	2,000	7 or more	38.7 %
Task 4	6,000	8 or more	19.4 %
Task 5	9,000	8 or more	19.4 %
Task 6	12,000	8 or more	19.4 %
Task 7	20,000	9 or more	7.3 %
Task 8	35,000	9 or more	7.3 %
Task 9	90,000	10 or more	1.9 %
Task 10	250,000	11 or more	0.3 %

For example, in Task 1, the selling price of security A exceeds the threshold if the price increases during 5 or more of the 12 months. This happens with probability 80.6%. The higher the threshold, the higher the number of required price increases. For example, in Task 2, the selling price exceeds the threshold if the price increases in 6 or more months. This happens only with probability 61.3%. Similarly, you can read in the subsequent lines how likely the threshold is met in the other tasks.

For comparison, recall that security B has a selling price of 17,000 Euros or 4,000 Euros, with equal probability, and a threshold of 10,000 Euros. Therefore, in each task, security B's selling price exceeds its threshold with probability 50%.

[end of insert for condition TREATMENT]

[page break]

Payment

For each round, one of the 10 tasks is picked at random. Each task is picked with equal probability by a computerized random draw. Depending on your decision in the task that is picked by the computer, you will receive the bonus or not.

After each round, you will learn the selling prices of both securities. We obtain these prices by means of computer simulation, which is conducted individually for each participant. You will receive the price information on a separate sheet of paper after each round. The selling price of your chosen security determines whether you receive the bonus in this round.

We then continue with another round of 10 tasks. (Recall there are 5 rounds.)

Are there questions about the tasks or payment rules in this experiment? If so, please raise your hand and we will help you at your desk.

If there are no further questions at this point, you will now face a brief understanding test. Only if you answer all questions correctly, you will proceed to the actual tasks.

In the top right corner of the understanding test, please enter the code number that you were assigned when you entered the laboratory. Please also enter this number on all subsequent sheets during this experiment.

[page break]

Understanding test

Code number: _____

Please record your code number on this sheet, as well as on all subsequent sheets during the experiment

Consider questions (1) to (4) below. You will only be allowed to continue with the experiment after answering all questions correctly. If you have a question of any kind, please raise your hand.

Questions:

- (1) Suppose you buy security B in the task that is picked by the computer. Suppose the selling price of security B is 17,000 Euros. Do you receive a bonus? _____
- (2) Suppose you buy security A in Task 1. Suppose the selling price of security A is higher than 100 Euros. Do you receive a bonus if Task 1 is picked by the computer?

(3) Suppose you buy security A in Task 10. Suppose the selling price of security A is less than 250,000 Euros. Do you receive a bonus if Task 10 is picked by the computer?

(4) Suppose you buy security A in both Task 1 and Task 10. Which of the two tasks has the higher chance that the selling price exceeds the threshold? _____

Once you finish the understanding test, please wait for instructions for the decisions. If you have a question, please raise your hand. Please make sure that the code number is recorded on the understanding test.

[page break]

[decision form]

Code number: _____

Round 1

	Threshold for security A:	Threshold for security B:	Your decision (A or B):
Task 1	100	10,000	
Task 2	500	10,000	
Task 3	2,000	10,000	
Task 4	6,000	10,000	
Task 5	9,000	10,000	
Task 6	12,000	10,000	
Task 7	20,000	10,000	
Task 8	35,000	10,000	
Task 9	90,000	10,000	
Task 10	250,000	10,000	

Once you finish making the decisions, please wait until the experimenter collects the decision sheets. If you have a question, please raise your hand. Please make sure that the code number is recorded on the first decision sheet.

[page break]

[feedback form]

Code number: _____

Selling prices in round 1:

Security A: _____

Security B: _____

[round 2 to 5 identically]

[page break]

SURVEY

Please provide the information requested below, but do **not** write your name. (Please respond truthfully to aid us in our research. You can be assured that all information will be stored in a 100% anonymous way, ensuring your privacy.)

CODE NUMBER _____

Date

_____.

Age: ____ **Sex:** ____ **Nationality:** _____

Undergraduate ____ **Graduate** ____ **Year of study** ____.

Main Subject of Study _____

Your average monthly budget, including all expenses for food and lodging:

Do you currently work for money? _____

Please indicate your main source of income: _____

In your household, do you live (check all that apply): ____ with parents ____ alone
____ with partner ____ with children ____ none of the aforementioned, but sharing an
apartment with someone else.

Did you take a mathematics course as an undergraduate? ____ yes ____ no

**Indicate the duration of schooling that your mother received, including any higher
education, by checking the number of years that comes closest:** ____4 ____8 ____12
____16 ____20

Indicate your father's years of schooling: ____4 ____8 ____12 ____16 ____20

**THE FOLLOWING ARE SOME NUMERICAL PROBLEMS. PLEASE ANSWER
THEM AS BEST YOU CAN.**

First problem: What is 15% of 1,000? _____

Second problem: A car rental agency charges \$35 a day plus \$0.14 per mile for its rental cars. If these charges include tax, what is the total cost of travelling 300 miles over 3 days in a car rented from this agency?

_____ \$42 _____ \$105 _____ \$125 _____ \$147 _____ \$300

Third problem: Which of the following is larger than $\frac{3}{5}$?

_____ $\frac{19}{35}$ _____ $\frac{13}{20}$ _____ $\frac{4}{7}$ _____ $\frac{7}{13}$ _____ None of the above

Fourth problem: If it takes 5 people 5 months to save a total of \$5,000, how many months would it take 100 people to save a total of \$100,000? _____

Fifth problem: A TV and a radio cost \$110 in total. The TV costs \$100 more than the radio. How much does the radio cost? _____

Sixth problem: In a lake, there is a patch of lily pads. Each day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake? _____