

Ensthaler, Ludwig; Nottmeyer, Olga; Weizsäcker, Georg; Zankiewicz, Christian

Working Paper

Hidden Skewness: On the Difficulty of Multiplicative Compounding under Random Shocks

DIW Discussion Papers, No. 1337

Provided in Cooperation with:

German Institute for Economic Research (DIW Berlin)

Suggested Citation: Ensthaler, Ludwig; Nottmeyer, Olga; Weizsäcker, Georg; Zankiewicz, Christian (2015) : Hidden Skewness: On the Difficulty of Multiplicative Compounding under Random Shocks, DIW Discussion Papers, No. 1337, Deutsches Institut für Wirtschaftsforschung (DIW), Berlin

This Version is available at:

<https://hdl.handle.net/10419/201410>

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.

1337

Discussion Papers

Deutsches Institut für Wirtschaftsforschung

2013

Hidden Skewness: On the Difficulty of Multiplicative Compounding under Random Shocks

(Updated Version August 2015)

Ludwig Ensthaler, Olga Nottmeyer, Georg Weizsäcker, and Christian Zankiewicz

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

IMPRESSUM

© DIW Berlin, 2013

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: On the difficulty of multiplicative compounding under random shocks^{*}

Ludwig Ensthaler[†] Olga Nottmeyer[‡] Georg Weizsäcker[§]
and Christian Zankiewicz[¶]

August 19, 2015

(Updated Version)

Abstract

Multiplicative growth processes that are subject to random shocks often have an asymmetric distribution of outcomes. In a series of incentivized laboratory experiments we show that a large majority of participants either strongly underestimate the asymmetry or ignore it completely. Participants misperceive the outcome distribution's spread to be too narrow-band and they estimate the median and the mode to lie too close to the distribution's center, failing to account for the compound nature of average growth. The observed biases are measured irrespective to risk preferences and they appear under a variety of conditions. The biases are largely consistent with a behavioral model in which geometric growth is confused with linear growth. This confusion is a possible driver of investors' difficulties with real-world financial products like leveraged ETFs and retirement savings plans.

JEL-Classification: C91, D03, D14, G02

Keywords: Behavioral Economics, Multiplicative Compounding, Skewness Neglect, Exponential Growth Bias

^{*}We thank Erik Eyster, Dorothea Kübler, Erik Mohlin, Peter Mörters, Tobias Schmidt, Adam Szeidl and Heinrich Weizsäcker and audiences at DICE, DIW Berlin, ESMT and the IEA World Congress 2014 for helpful comments and we thank colleagues at the decision laboratories of Technical University Berlin and University College London, especially Brian Wallace and Mark Henninger, for their excellent contributions in the preparation and conduct of the experiments. The paper has much benefitted from the editorial work of the Editor, an Associate Editor and three referees. Financial support by the ERC (Starting Grant 263412) is gratefully acknowledged.

[†]Rocket Internet and WZB Berlin Social Science Center, ludwig@ensthaler.com

[‡]IZA Bonn, nottmeyer@iza.org

[§]Humboldt University Berlin and DIW Berlin, weizsaecker@hu-berlin.de

[¶]DIW Berlin, czankiewicz@diw.de

1 Introduction

Many household investors face a particular mismatch in the time frames of asset return evaluations. They acquire their most important financial assets with the intention to liquidate them in the relatively distant future but the available return information concerns much shorter time intervals. Real estate investments, retirement savings plans or investments in college funds share this feature. In all of them, the relevant outcomes are the investments' performances over several decades but the available information concerns their short-term performances, like 1-year returns. To forecast the return on the planned (or any plausible) distant selling date, an investor needs to extract the price distribution at the selling date by compounding the available short-term return distributions. This is a formidable task for the average person.

Two possibly severe biases may arise when forecasting the distribution of long-run growth. First, one may fail to compound the effects of multi-period growth. Second, one may ignore the skewness that arises over time and, e.g., confuse the mean return with the median return. The first of these biases can also appear in deterministic settings—there it is established as “exponential growth bias”. When asked to assess the total effect of accumulating 7% growth for ten periods, a substantial fraction of respondents give an answer that is closer to 70% than to the actual 97%. The analyses of Stango and Zinman (2009) and Levy and Tasoff (2014) indicate that the bias is empirically relevant as it affects households' borrowing and saving decisions.¹ The second bias—which we call “skewness neglect” hereafter—is less well known in the academic literature² but among investment practitioners and financial market regulators there are related discussions of it. Investors apparently need to be made aware that the compounding of random growth can transform a symmetric 1-period return distribution into a skewed multi-period return distribution. An important real-world example of this is the family of leveraged exchange-traded funds (leveraged ETFs). These assets have a highly volatile and fairly symmetric 1-period volatility; holding them for multiple periods results in severe skew.

This paper presents a series of incentivized laboratory experiments that extend the research on the perception of multiplicative growth to the stochastic domain and accounts for both of the above-described biases. As an example that demonstrates skewness neglect, consider the stylized experiment of Section 3.3. A very volatile asset either increases in value by 70% or decreases in value by 60% in every period, each growth rate realizing with a chance of one half. If the investor buys the asset she must hold it for twelve periods. With an initial value of 10,000, what would the asset likely be worth at the end of period 12? To ask this question in an incentive-compatible way, we let the participants bet on five possible outcome ranges for the period-12 value of the asset: a) up to 6,400, b) between 6,400 and 12,800, c) between 12,800 and 19,200, d) between 19,200 and 25,600, or e) above 25,600.

¹Both Stango and Zinman (2009) and Levy and Tasoff (2014) present survey evidence of a statistical connection between the bias and respondents' savings behaviors. Levy and Tasoff (2014) also analyse theoretical implications of the bias, e.g., an overestimation of future income that arises from too-moderate time discounting of income. The effect can result in overconsumption if income is shifted to later time periods. Related effects are addressed in the experiments by McKenzie and Liersch (2011).

²The only other academic study that we are aware of is by Stutzer and Jung Grant (2010), discussed in the next Section.

We then simulate the random process and if the simulated path ends up in the outcome range that a participant has bet on, she receives a bonus of 20 Euro. If not, she receives nothing. The most popular answer is c), chosen by 43% of participants, followed by d) (28%) and b) (17%). Response options a) and e) come tied bottom with a mere 6% of responses each. However, the optimal response is a); the median of the resulting distribution is 989 and the probability that the process ends up in the lowest interval is 80%. A simple reasoning for this is that a value increase of 70% cannot recover a value decrease by 60%, hence most trajectories have a downward trend and the distribution is highly skewed already after 12 periods. The participants fail to realize this and instead report answers that are consistent with a confusion of mean and median. Their average payoff (based on their decisions) amounts to a meager 2 Euro in this experiment, whereas the optimal response would earn them 16 Euro in expectation.

Our series of experiments examines this kind of mistake systematically, finding that participants' perception of stochastic growth deviates in predictable ways from the rational prediction. Both of the above biases are found to be relevant. Overall, the experimental results are in line with a simple model of misperception of compounding of shocks. This model, which we label "linearity bias model", formulates the hypothesis that the agent fails to calculate the effects of compounding and that this failure appears not only in deterministic but also in stochastic settings. The agent perceives a linear evolution in the sense that she perceives the distributions of absolute changes as constant over time, instead of the relative changes being constant over time.³ In effect, all multiplicative growth is mistaken as additive growth with a constant distribution of increments. The model thereby predicts both exponential growth bias and skewness neglect.

Importantly, the model also allow predictions about the strength of the two effects. It predicts that the agent has a fairly rational perception of the growth process in the case that both per-period volatility and per-period return are low. For larger volatility, skewness neglect becomes relevant and will lead to an overestimation of the median. Increasing the per-period return to a sufficient degree makes the exponential growth bias dominate (no matter what the volatility), such that the agent underestimates the median. This somewhat intricate pattern of predictions cannot be generated by any of the biases alone but it is confirmed by the experimental data. Section 3.2 shows these effects in our main design, where we simply ask the participants to predict the most likely outcome of a growth process. In the binomial-tree processes that we use, the most likely outcome is also the median and thus the responses can be used for assessing the subjectively perceived medians. All of our experiments are incentivised, in ways that make truth-telling optimal irrespective of one's risk attitudes.

The experiments of Sections 3.3, 3.4, 4 and 5 test the predictions of the model in different variations of the experimental setting. We vary the incentive schemes, the level of feedback as well as the nature of both the underlying asset and investment strategy that the participants are asked to assess. Thereby, we can investigate the robustness of the effects and we can also inquire about several other implications of the linearity bias model. Qualitatively, almost all predictions of the model are borne out in

³The model is akin to that of Levy and Tasoff (2014) albeit developed independently.

the data, and often with large discrepancies to the rational prediction. For example, in treatments with a considerably high return volatility, about ninety percent of participants overestimate the median. The model is also fairly successful in predicting the participants' misperceptions of the 10th and 90th percentile of the relevant long-run distributions: the 90-10 spread is generally underestimated. This holds both for binomial-tree assets and for more realistic assets that we base on the historical returns of the German DAX index.

The rest of this paper is organized as follows. Section 2 briefly discusses related literatures. Section 3 (Study 1) introduces the linearity bias model and the main experimental design. In Subsection 3.2 we report four experimental treatments that investigate the perceived medians and produces the above-described data pattern of exponential growth bias and skewness neglect. In Subsections 3.3 and 3.4 we vary the elicitation method, validating the effect of skewness neglect. Section 4 describes Study 2, which varies the format of Subsection 3.4 to investigate the perceived 10th and 90th percentiles. This study also varies the investment horizon. Section 5 (Study 3) reports on the extension to a setting with a real-world asset based on the German DAX index. Section 6 concludes and the Appendix contains further experimental treatments.

2 Review of related literature

Classic studies in cognitive psychology discuss quite extensively to what degree the human cognitive apparatus is able to account for the distinction of linear versus nonlinear relations between variables. Wagenaar and Sagaria (1975) ask participants to predict an exponential data series representing an index for pollution. They find that participants strongly underestimate exponential growth. Wagenaar and Sagaria (1978) show that underestimation of exponential growth is robust to the amount of information available to the participants and Wagenaar and Sagaria (1979) show that the effect is robust to the framing of the information. Kemp (1984) surveys perceptions of changes in the cost of living. Respondents systematically underestimate the increase in cost, which is also in line with a misperception of exponential growth. Much of the early data analysis uses responses to quiz-type questions, but a subsequent specialization of this literature more and more focuses on economic contexts, like the perception of compound growth from interest or loan payments. Eisenstein and Hoch (2005), Stango and Zinman (2009), Christandl and Fetchenhauer (2009), McKenzie and Liersch (2011) and Levy and Tasoff (2014) document that participants underappreciate the effects of compound interest and thereby predictably underestimate the compound effect of growth. Chen and Rao (2007) show that retailers can strategically use this bias by posting double dip price discounts (a discount of 20% followed by another 25% discount is perceived to be a 45% reduction, not the actual 40%). As described in the Introduction, our paper can be viewed as an extension of this literature to non-deterministic growth processes.

An important predecessor of our paper is the study by Benartzi and Thaler (1999) who, among other things, study biases in the compounding of the long term distributions from a given short term dis-

tribution. Their experimental participants choose different hypothetical retirement plans depending on whether they are given the historical return distribution of retirement plans for a 1-year period or a 30-year period. Benartzi and Thaler (1999) relate this bias to the effects of myopic loss aversion (see also Samuelson (1963), Redelmeier and Tversky (1992), Gneezy and Potters (1997), and Klos, Weber, and Weber (2005)). While we agree that myopic loss aversion likely plays a role in households' long term investment decisions, our experiments suggest that household decisions can also be misguided by a biased perception of the underlying growth processes.⁴ This is also consistent with the only experimental paper on skewness neglect that we found, by Stutzer and Jung Grant (2010). Their hypothetical investment experiments find an inflated investment rate in treatments where their participants have to calculate the compound return by themselves.⁵

Another related literature studies whether experimental participants have a correct understanding of financial options. We refer the reader to Gneezy (1996) and Abbink and Rockenbach (2006) for previous results in this—surprisingly small—literature. We note that the assets that we use in Sections 3 and 4 have the same structure as the underlying asset in Cox, Ross, and Rubinstein's (1979) well-known model of European call options. A consistent finding of misperceptions of such assets may therefore indicate a potential mispricing. This is not further studied in our paper, which focuses on investments in the underlying asset itself.

3 Study 1: Assessments of Median and Mode

Study 1 is composed of the three partial studies 1(a), 1(b), and 1(c), each testing participants' perceptions of an asset's mode or median. We consider one of the most elementary asset covered in the finance literature: the binomial-tree asset with fixed maturity (Cox, Ross and Rubinstein, 1979). This asset's multiplicative growth μ_t is a binary random variable with a constant 50-50 chance of moving up or down for T periods, i.e. $\mu_t \in \{\mu^h, \mu^l\}$ where the percental uptick $\mu^h \geq 0$ and the percental downtick $\mu^l \geq 0$ are equiprobable in each $t = 1 \dots T$.

3.1 The Linearity Bias

We start the analysis by presenting our simple model of biased decision making. Consider a decision maker who ignores compounding of interest: when asked to predict the accumulated value gain of an investment that yields a per-period interest of r over T periods, she quotes a total gain of rT . That is, she wrongly perceives the absolute changes, not the relative changes, to be constant across

⁴A distinction between our study and the existing experimental work on myopic loss aversion is that the existing papers largely make use of additive growth processes.

⁵The experiment by Stutzer and Jung Grant (2010) uses a quite similar experimental wording as the experiment described in Section 3.4 and in our paper's first (2010) version despite having been developed and written independently. A separate and important experimental literature examines the preferences regarding skewness, see Deck and Schlesinger (2010), Br  nner, Levinsky and Qiu (2011), Ebert and Wiesen (2011), and Eckel and Grossman (2014). We restrict this paper to the perception of the distribution, not its valuation.

the periods. This feature is our model’s sole bias—linearization bias hereafter—and we can readily extend it to the domain of stochastic growth.

Formally, let Y_0 denote the known initial price of an asset with a random price series $\{Y_0, Y_1, \dots\}$ and let η_t be the random variable describing the absolute price growth in t , e.g. $Y_1 = Y_0 + \eta_1$. The linearization biased (LB) decision maker views the distribution of η_t as constant across t , with a distribution identical to that of η_1 . As a result, the LB decision maker misses out on all effects of multiplicative compounding, which may or may not occur in the true growth process. If multiplicative growth occurs, η_t ’s distribution changes across t but the LB decision maker ignores it. The absolute change in $t = 1$ has a special role in that it is assumed to be readily available to the agent. This assumption reflects our aim of capturing the neglect of compound interest, as $t = 1$ is the only period where η_t is known without compounding.⁶

To investigate the effects of the bias, Study 1 and Study 2 consider binomial-tree price series where the relative, not absolute, price growth is constant in t by definition. Let μ_t be the random variable describing the relative price changes occurring in t , e.g. $Y_1 = Y_0\mu_1$. Binomial-tree models, like many other natural growth processes, have the distribution of μ_t i.i.d. across t . An unbiased decision maker perceives the true distribution of the period- T price as $Y_T = Y_0 \prod_{t=1}^T \mu_t$, with $\{\mu_t\}$ i.i.d., whereas the LB decision maker perceives the final price as $\tilde{Y}_T = Y_0 + \sum_{t=1}^T \tilde{\eta}_t$, with $\{\tilde{\eta}_t\}$ i.i.d. and its distribution equal to that of η_1 .

It is easy to see that for binomial-tree models with equiprobable upticks and downticks, the LB model predicts full skewness neglect: the percentiles of the perceived \tilde{Y}_T are located in a symmetric way around \tilde{Y}_T ’s median because symmetry is preserved under addition of random variables. With a strictly positive per-period average growth, the model also predicts the exponential growth bias: the final distribution’s mean is underestimated, $\mathbb{E}[\tilde{Y}_T] < \mathbb{E}[Y_T]$.

In the following, we use the LB model to derive qualitative (directed) predictions of biased decision making in our contexts, generated by simple numerical applications.⁷ Our main empirical focus lies on measuring the perception of the median of Y_T . Under a rational perception of Y_T ’s distribution, median and mean differ in binomial trees. LB decision makers, in contrast, perceive them as identical. Another relevant property is that the median is identical to the mode of Y_T ’s distribution, to which both optimal and LB decision makers agree (despite disagreeing on the value). This allows us to formulate alternative elicitation tasks, equivalently asking for median or mode.

⁶The assumption is also consistent with the evidence on the exponential growth bias in deterministic settings: it prescribes that the slope of linearly perceived growth is given by the initial slope. Formally, assume a constant multiplicative growth path such that $Y_T = Y_0(1 + r)^T$, where r is interest. The LB decision maker perceives a constant additive growth instead, $\tilde{Y}_T = Y_0 + T\tilde{\eta}$ where $\tilde{\eta} = Y_1 - Y_0 = rY_1$. For $r > 0$ and $T > 1$, we have $\tilde{Y}_T < Y_T$.

⁷We also discuss the model’s point predictions for completeness; but as a model of such simplicity cannot plausibly capture the precise decision process we focus the statistical analysis on the qualitative predictions.

3.2 Study 1(a): Biased perception of the mode

3.2.1 Experimental Design

Participants in Study 1(a) are presented with a security whose price is currently at $Y_0 = 100$ and changes by a factor $\mu \in \{\mu^h, \mu^l\}$ with equal probabilities during each period and with all random draws being independent. The participants' task is to locate the mode of the security's outcome distribution after $T = 12$ periods.⁸ The task is made incentive compatible as follows. After a participant's response, the experimenter simulates a set of 100 values of Y_T . If at least one of these simulated values differs by less than 1 from the participant's stated value she receives a bonus of 20 Euro, otherwise not. The procedure thus prompts the participant to report the location (more precisely, an interval of length 2) where she perceives Y_T 's highest likelihood. The optimal response would be to report the mode of Y_T . Notice that reporting the mode is optimal irrespective of risk preferences and of the nature of the perceived Y_T : the incentive scheme uses only two possible payments—receive a bonus versus not—making it optimal for any participant with monotonic preferences to maximize the subjectively perceived probability of receiving the bonus by stating the price that she thinks is most likely. The procedure is also simple to understand and allows asking a straightforward question about the participants' prediction of the asset's price evolution.⁹

There are no repetitions in Study 1 and thus participants only give one statement per security. Each participant is, however, asked to report a prediction on two different securities, to increase the number of observations and to allow for within-person comparisons. One of the two responses is randomly picked to be payoff relevant, at the experiment's conclusion.

Overall, Study 1(a) covers four different securities that differ in the values $\{\mu^h, \mu^l\}$, appearing pairwise in two treatments. Participants in treatment 1 assess Security 1's ($\mu^h = 1.7, \mu^l = 0.4$) and Security 2's ($\mu^h = 1.075, \mu^l = 1.025$) modes, whereas participants in treatment 2 assess Security 3's ($\mu^h = 1.4, \mu^l = 0.7$) and Security 4's ($\mu^h = 1.8, \mu^l = 1.1$) modes. Each participant thus faces one security which can depreciate as well as appreciate, and one security which can only appreciate. In treatment 1, the two securities have identical means but different per-period volatilities (as measured by the spread $(\mu^h - \mu^l)$) and in treatment 2, the two securities have different means but identical per-period volatilities. Moreover, the mean of Security 3 is identical to that of Securities 1 and 2. Participants are randomly assigned to treatments 1 or 2. To account for possible learning effects, the order of the two securities randomly varies between the participants within a treatment.

All 127 participants (63 in treatment 1 and 64 in treatment 2) are students at Technical University Berlin. Six sessions, three in each treatment, are conducted in a computer-based format using the software z-Tree (Fischbacher, 2007). Participants receive a participation fee of 5 Euro in addition to their possible bonus of 20 Euro.¹⁰

⁸The decision problems are phrased in a financial investment context. Its descriptions begins with the wording: "You are a manager and have to make a decision..."

⁹All instructions are in an online appendix.

¹⁰Before the experiment starts, participants also face an understanding test which they have to answer correctly before they may proceed. Questions are carefully chosen to not suggest any responses to the participants.

3.2.2 Linearity Bias Prediction

The securities in Study 1(a) are specified such that they allow predictions about the relative strengths of the above-described two effects, exponential growth bias and skewness neglect. To examine the effect of a large per-period volatility on our participants' perceptions, we consider Security 1 ($\mu^h = 1.7, \mu^l = 0.4$) with -60% and +70% as possible percentage changes and predict a strong effect of skewness neglect. With a perceived constant distribution of absolute changes $\tilde{\eta}_t \in \{-60; +70\}$, the LB decision maker perceives a symmetric distribution of the period-12 selling price of this security with mode, mean, and median at $\mathbb{E}(\tilde{Y}_T) = Y_0 + \mathbb{E}(\sum_{t=1}^T \tilde{\eta}_t) = 160$. Thus, skewness neglect leads to a strong overestimation of Security 1's true mode at 9.89. The exponential growth bias by itself makes no biased prediction for Security 1.

For a relatively lower per-period volatility, as captured by Security 3 ($\mu^h = 1.4, \mu^l = 0.7$) with -30% and +40% as possible percentage changes, skewness neglect becomes less extreme and results in a weaker, but still sizable, overestimation of the median: the LB model predicts a response of 160 instead of the correct 88.58.

Further decreasing per-period volatility, as for Security 2 ($\mu^h = 1.075, \mu^l = 1.025$), results in the predictions that the LB decision maker has a fairly rational perception of the growth process: she perceives the period-12 price of Security 2 at 160 instead of the correct 178.97. That is, even a model allowing for both skewness neglect and the exponential growth bias is fairly ineffective and predicts that Security 2's price is only mildly underestimated.

Security 4 ($\mu^h = 1.8, \mu^l = 1.1$) allows a strong effect in the opposite direction, which is due to the exponential growth bias. Here, the LB decision maker perceives the most likely period-12 price at 640, while the true mode of the price distribution lies at 6,025.47. With such high per-period mean growth, the question of symmetric versus asymmetric percentiles (skewness neglect) becomes less important than the effect of the exponential growth bias.

3.2.3 Results

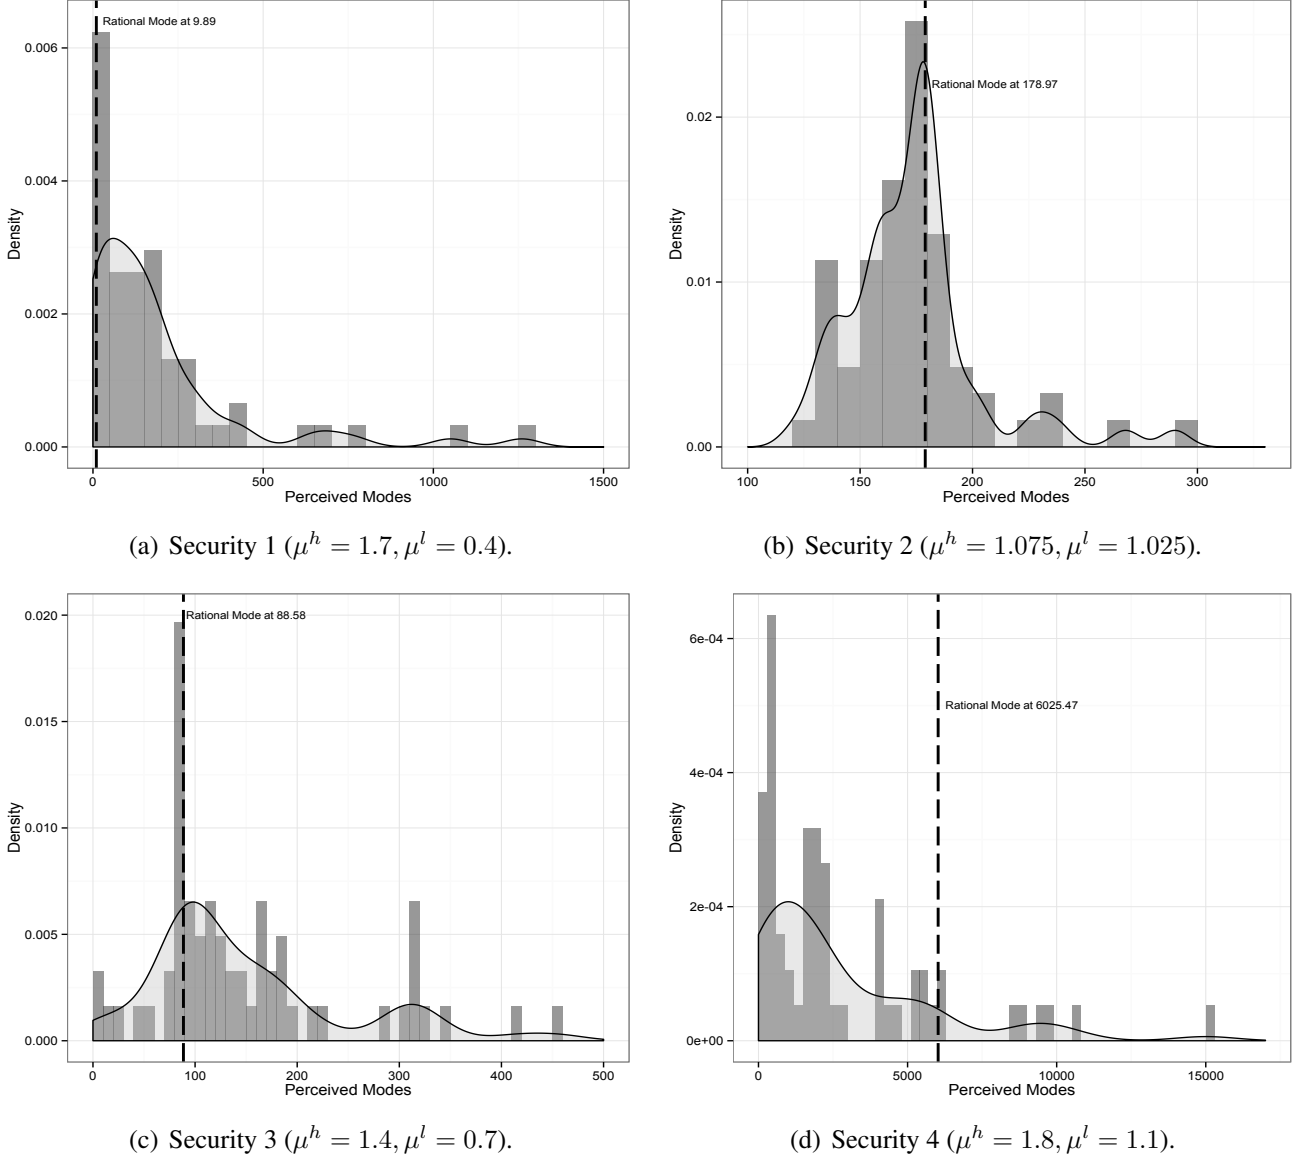


Figure 1: Densities of the subjectively perceived modes for Securities 1 through 4. Dashed lines indicate rational benchmarks.

Figure 1 illustrates the distributions of subjective mode perceptions for all four securities. The participants' predictions for Security 1 ($\mu^h = 1.7, \mu^l = 0.4$) are in the figure's panel 1(a) and show a substantial degree of overestimation. Consistent with skewness neglect, 87 % of our participants overestimate the true mode. The frequency of overestimation lies significantly above 50% (p-value<0.001, one-sided binomial test). Although the data show a peak in the neighborhood around the optimal value, most participants' degree of overestimation is substantial. Half of them predict the distribution's mode to lie above 120—more than 12 times the true value.

Figure 1(c) illustrates the subjective mode perceptions for Security 3 ($\mu^h = 1.4, \mu^l = 0.7$). As for Se-

curity 1, the data show a notable proportion of participants, 70%, overestimating the mode. While the frequency of overestimation lies significantly below that of Security 1 (p-value<0.001, two-sample binomial test), it is still significantly greater than 50% (p-value<0.001, one-sided binomial test). These observations are consistent with the LB model in the sense that the model predicts an overestimation for both securities and a larger overestimation for Security 1 than for Security 3. But also apart from the model, the comparison between Security 1 and Security 3 is relevant as it shows the effect of skewness neglect in isolation: the mean is constant between them whereas the higher volatility in Security 1 changes median and mode. The higher proportion of overestimation in Security 1 shows that participants do not fully appreciate this difference.

The participants' assessments of Security 2 ($\mu^h = 1.075$, $\mu^l = 1.025$) are depicted in Figure 1(b) and (again consistent with the LB model) show a different picture. Only a minor underestimation of Security 2 appears: 57% of participants state mode values below the true mode of 178.97, a proportion that is not significantly greater than 50% (p-value>0.15, one-sided binomial test). Moreover, the median response is not significantly different from the optimal value (p-value>0.05, Wilcoxon signed-rank test). Neither skewness neglect nor the exponential growth bias show to be relevant for this security. The perceptions for Security 4 ($\mu^h = 1.8$, $\mu^l = 1.1$), with a higher average per-period return, are illustrated in Figure 1(d) and show a substantial degree of underestimation. 89% of the participants state responses that lie below the true mode. This share is significantly larger than 50% (p-value<0.001, one-sided binomial test) and also significantly larger than the share of participants who underestimate Security 2's mode value (p-value<0.001, two-sample binomial test). Once again, these observations are consistent with the much stronger prediction of the LB model for Security 4 than for Security 2. Moreover, it is notable that the data confirm the LB model's prediction that the exponential growth is more relevant than skewness neglect in Security 4.

Security 4 is also well suited to show that participants do not simply confuse median and mean. Stating the security's mean value (8,638.06) instead of the median / mode would have resulted in an overestimation, in contrast to the LB model's prediction and to the evidence in Figure 1(d).

3.3 Study 1(b), Robustness I: Coarser choice

Study 1(b) elicits the modes of our participants' subjective distributions using a coarse choice list design with steep incentives. The task description is also shorter and simpler than in Study 1(a). (This is the simple design described in the paper's Introduction.)

3.3.1 Experimental Design and Linearity Bias Prediction

The experimental asset of interest, Security A, follows a binomial-tree +70%/-60% process over 12 periods that is identical to Security 1 in Study 1(a) with the sole exception that its initial price is 10,000. In three experimental sessions, 69 students at Technical University Berlin are presented with the growth process of Security A and are asked to pick one out of five investment opportunities, la-

beled Investment 1 through Investment 5, whose return depends on the period-12 price of Security A. Just as in Study 1(a), Study 1(b) also ensures incentive compatibility under a wide set of preferences by using only two possible payments per choice problem—receive a bonus of 20 Euro versus not. Participants are told that Investment 1 “makes a gain” (in effect pays the bonus—see below) iff the selling price of Security A is between 0 and 6,400. Investment 2 makes a gain iff the selling price of Security A is between 6,400 and 12,800, Investment 3 makes a gain iff the selling price is between 12,800 and 19,200, Investment 4 makes a gain iff the selling price is between 19,200 and 25,600, and Investment 5 makes a gain iff the selling price is above 25,600.

Once a participant has chosen his or her investment the computer simulates the period-12 price of Security A, with a separate simulation for each participant. If the chosen investment makes a gain, the participant receives the bonus. Under any belief the decision-maker should, evidently, choose the interval that has the largest probability of containing the period-12 price. Due to the price distribution’s large skew, the rational prediction is to choose Investment 1, whose corresponding interval $[0, 6400]$ contains the period-12 price with 80% chance. The other intervals thus have a far lower success chance and the monetary incentive for a participant to choose one of them is far lower. An LB decision-maker, however, perceives a symmetric price distribution \tilde{Y}_T with a mode at 16,000. This implies that the interval containing 16,000 has the highest perceived chance of yielding the bonus. An LB decision maker therefore chooses Investment 3.

3.3.2 Results

The numbers of participants (and percentages in parentheses) choosing Investments 1 through 5 are $\{Inv. 1 : 4 (6\%), Inv. 2 : 19 (28\%), Inv. 3 : 30 (43\%), Inv. 4 : 12 (17\%), Inv. 5 : 4 (6\%)\}$. The distribution is significantly different from uniform choice (p-value<0.001, chi-square test) and indicates no tendency to choose a mode near zero. Overall, with 94% of our participants significantly more than half of the sample overestimate the true mode (p-value<0.001, one-sided binomial test). While only 6% make the optimal choice of Investment 1, 43% conform with the linearity bias model and choose Investment 3. The participants give up significant amounts of money due to the bias: While the optimal choice would earn 16.12 Euro in expectation, the observed choice distribution on average earns only 2.07 Euro in expectation per participant.

3.4 Study 1(c), Robustness II: Choice list format and repetitions

In Study 1(c) we use a choice list mechanism to identify bounds on the median of each participant’s subjectively expected distribution. We also let our participants repeat this task over five rounds.¹¹

¹¹In a further treatment variation of Study 1(c), we additionally provide the participants with an explicit calculation of the distribution of compound price changes for the respective security and we point out the asymmetry in the price distribution. The observed choice bias decreases strongly in this treatment, consistent with the presumption that the bias stems from a cognitive problem and is not driven by the particular choice format. A detailed description of this treatment is in Appendix A.

3.4.1 Experimental Design

The experimental responses are bounds on the participants' subjectively perceived median of returns, irrespective of risk preferences. As for the two previous studies, Study 1(c) ensures incentive compatibility under a wide set of preferences by using only two possible payments per choice problem. In each round of the experiment, two risky securities are on offer and the selling price of the chosen security determines whether or not the participant receives the bonus. Security A is identical to the asset in Study 1(b), with a +70%/-60% price change in each of the 12 periods and an initial price of 10,000. 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 experimental design we describe Security B analogously to Security A, with the difference that Security B has only a single price equiprobable change of +70% or -60% 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 throughout the experiment (hence Security B holds a 50-50 chance of receiving the bonus) whereas the threshold t_A varies between 10 different values. 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 distribution of Security A. Table 1 lists the 10 choice problems (Task 1, Task 2, etc.) as seen by the participants. Given that the true median of Security A's selling price is 989, the rational prediction is for the participants to choose A in Task 1 and Task 2 and to choose B in all subsequent tasks.

	Thresholds for Security A	Thresholds for Security B	Your decision (A or B)
<i>Task 1</i>	100	10,000	—
<i>Task 2</i>	500	10,000	—
<i>Task 3</i>	2,000	10,000	—
<i>Task 4</i>	6,000	10,000	—
<i>Task 5</i>	9,000	10,000	—
<i>Task 6</i>	12,000	10,000	—
<i>Task 7</i>	20,000	10,000	—
<i>Task 8</i>	35,000	10,000	—
<i>Task 9</i>	90,000	10,000	—
<i>Task 10</i>	250,000	10,000	—

Table 1: The 10 binary choices.

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.¹²

Three sessions are conducted in a paper-and-pencil format, with 68 student participants at Technical University Berlin. 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 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.

3.4.2 Linearity Bias Prediction

Analogous to the previous studies, the LB model predicts that participants overestimate the true median (989) by an order of magnitude and chooses a switching value in the interval $[12,000 - 20,000)$. For notation, let $q_{0.5,i}$ be the elicited *lower* bound of participant i 's assessment of the median: i invests in Security A for all values $t_A \leq q_{0.5,i}$ and invests in Security B for all strictly larger t_A . For the sake of simplicity we restrict attention to cases where participants' choices reveal such a unique switching value, a property that is true in 93% of our data.¹³ By analogy, let $q_{0.5}$ be the rational benchmark for $q_{0.5,i}$ (dropping the subscript i), i.e. the lower bound of the median that would be elicited from a decision maker who behaves optimally. Here and elsewhere in the paper, we focus on revealed lower bounds when applicable.

3.4.3 Results

Range of subjective median	Share of participants switching from A to B				
	Round 1	Round 2	Round 3	Round 4	Round 5
[0 – 100)	0.018	0.000	0.000	0.018	0.000
[100 – 500)	0.000	0.000	0.000	0.000	0.035
[500 – 2,000)	0.000	0.054	0.072	0.072	0.107
[2,000 – 6,000)	0.036	0.145	0.127	0.200	0.303
[6,000 – 9,000)	0.107	0.090	0.254	0.309	0.142
[9,000 – 12,000)	0.411	0.381	0.309	0.236	0.196
[12,000 – 20,000)	0.196	0.181	0.109	0.127	0.142
[20,000 – 35,000)	0.179	0.090	0.109	0.036	0.053
[35,000 – 90,000)	0.054	0.054	0.000	0.000	0.017
[90,000 – 250,000)	0.000	0.000	0.000	0.000	0.000
[250,000 – ∞)	0.000	0.000	0.018	0.000	0.000

Table 2: Subjective median ranges over the five rounds.

¹²Each additional round comes with the chance to earn a new bonus 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 , under a wide set of preferences for choice under uncertainty.

¹³If a participant has multiple switching points in one round, her answers in the remaining rounds are still considered in our data analysis. None of our conclusions would change if we dropped all responses by participants 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.

Table 2 lists the implied ranges for the medians of the participants' subjective distributions of Security A's selling price over the five rounds. In round 1, not a single participant reveals a subjective median between 500 and 2,000 (i.e., optimal switching at Task 3). Instead, with 98% of participants significantly more than 50% (p-value<0.001, one-sided binomial test) reveal that their subjective medians are above 2,000. The results of rounds 2 to 5 show that a significant proportion of 86% of participants overestimate the median still in round 5 (p-value<0.001, one-sided binomial test for >50%). The modal choice in round 1 (41% of participants) indicates a subjective median between 9,000 and 12,000, with the next-higher interval [12,000-20,000) attracting 20% of participants' choices.

4 Study 2: Other Quantiles and Other Growth Processes

In this section we examine the robustness of the linearity bias predictions with respect to variations of the investment horizon and the choice of quantiles of the participants' subjective distributions that we elicit. The results confirm the LB model's implication that the compound distribution is perceived as too symmetric and too narrow-band if there is substantial randomness in the growth process.

Subsection 4.1 describes the experimental design in detail. Subsection 4.2 contains the prediction of the LB model and Subsection 4.3 has the results.

4.1 Experimental Design

There are two treatments in Study 2, both involving assets similar to the processes described in Study 1. Participants can buy Security A at a price of 100. If they buy it they have to sell it after T^k periods, where k indexes the treatment. The price moves by about 20 percent in each period: In both treatments, *High Volatility_Short* (HVS) and *High Volatility_Long* (HVL), the parameters specifying upticks and downticks are $\mu^{h,HVS} = \mu^{h,HVL} = 1.212$ and $\mu^{l,HVS} = \mu^{l,HVL} = 0.811$. The sole difference between these two treatments is in the length of time until maturity: $T^{HVS} = 14$ and $T^{HVL} = 140$.¹⁴

As in Study 1(c), a participant of Study 2 who buys Security A receives a fixed bonus if the selling price at maturity exceeds a given threshold t_A . These thresholds differ between treatments and are listed in Table 3. The alternative choice option is Security B which yields the bonus with a certain probability.¹⁵ To elicit three different quantiles about Security A's selling price, Security B has three different specifications. Each participant faces each specification once. Security B1 yields the bonus

¹⁴In two further treatments *Low Volatility_Short* (LVS) and *Low Volatility_Long* (LVL), with $\mu^{h,LVS} = \mu^{h,LVL} = 1.012$ and $\mu^{l,LVS} = \mu^{l,LVL} = 1.011$, the price motion is approximately deterministic (i.e., the price volatility is very low) and the price has positive growth with certainty. The number of time periods until maturity is analogous to HVS and HVL, at $T^{LVS} = 14$ and $T^{LVL} = 140$. A further treatment *Low Volatility_Long_NoCalculator* (LVLNC) is identical to LVL but does not grant the participants access to calculators. The results can be found in Appendix B.

¹⁵Different from Study 1(c), the instructions simply report to the participants the probability with which Security B yields the bonus, without referring to a separate threshold t_B .

with 90% probability, B2 with 50% and B3 with 10%. Accordingly, each participant faces three choice lists. First, she chooses between Securities A and B1 for the different thresholds of Security A. This allows us to infer bounds on her subjective 10th percentile of Security A's selling price. For example, suppose that participant i in treatment HVS chooses Security A over Security B1 in Task 1 and Task 2 and chooses Security B1 over Security A in tasks 3 through 10. Inspecting Table 3 (first column) we see that this is subjectively optimal iff participant i 's subjective 10th percentile for Security A's selling price is between 30 and 45. In line with our previous notation, we would thus record the elicited lower bound of i 's subjective quantile for Security A's selling price as $q_{0.1,i}^{HVS} = 30$. As her second set of tasks the participant faces the analogous choices between Securities A and B2 (with the same list of thresholds for Security A). This allows us to infer a lower bound on her subjective median of the selling price, $q_{0.5,i}^{HVS}$. Finally, she faces the analogous list of choices between Security A and B3, allowing us to infer a lower bound on the 90th percentile of the same price, $q_{0.9,i}^{HVS}$.¹⁶

	Values of t_A <i>HVS</i>	Values of t_A <i>HVL</i>
<i>Task 1</i>	15	2
<i>Task 2</i>	30	5
<i>Task 3</i>	45	15
<i>Task 4</i>	65	60
<i>Task 5</i>	95	140
<i>Task 6</i>	125	230
<i>Task 7</i>	155	350
<i>Task 8</i>	190	550
<i>Task 9</i>	225	700
<i>Task 10</i>	265	1,000

Table 3: The thresholds t_A by treatment condition.

The computer terminals report feedback to the participants in the form of a sample selling price of Security A. In each treatment, this concludes the first round of the experiment. The experiment is then repeated for four additional rounds. Hence, the experiment comprises five identical rounds for each participant.

All 58 participants are undergraduate students at University College London. Nine sessions are conducted in a computer-based format using the software z-Tree (Fischbacher, 2007). Each of the two treatment conditions is faced by a random subset of participants in each session, without making them aware that other participants face different treatment conditions (N=29 in both, HVS and HVL).

In contrast to Study 1, participants in these sessions are supplied with a hand-held calculator that they can use throughout the experiment. The protocol is fixed across all sessions. First, printed instruc-

¹⁶After the elicitation of the subjective quantiles we also ask for the participants' beliefs of Security A making a profit. We do not use the resulting data in the analysis but refer to the paper's previous version (Ensthaler et al., 2013) and to the instructions for a description of the experimental details and the results.

tions are distributed and participants have to pass an understanding test.¹⁷ Then the computer-based experiment commences and guides the participants through the five rounds in immediate succession. For a simpler data analysis, the participants' computer interfaces restrict responses to satisfy two constraints. First, responses must exhibit at most one switching point on a choice list between Security A and a single B-type security. That is, a participant cannot switch back and forth between Security A and the respective B-type security. Second, the elicited quantiles must be ordered in a consistent way: participants cannot switch from Security A to B1 at a threshold that exceeds the threshold at which she switches from Security A to B2, which in turn cannot exceed the threshold at which she switches from Security A to B3.¹⁸

Participants receive a participation fee of £5 and a possible bonus of £5 per round. In each round, a single choice is randomly determined to be payoff-relevant, giving an ex-ante incentive to act optimally in each task.¹⁹

4.2 Linearity Bias Prediction

The treatment comparison in Study 2 focuses on the effects of an increased time horizon. With a perceived constant distribution of absolute changes $\tilde{\eta}_t \in \{-18.9; 21.2\}$, the LB decision maker perceives selling price distributions with mode, mean, and median at 116.10 in treatment HVS and 261.00 in treatment HVL. The LB model thus predicts that participants overestimate the true median (88.64 in HVS and 29.96 in HVL), mildly for the former treatment and much stronger for the latter. (The LB model's lower bound point predictions are $q_{0.5,i}^{HVS} = 95$ in treatment HVS with the rational assessment at 65, and $q_{0.5,i}^{HVL} = 230$ in treatment HVL with the rational assessment at 15.)

Moreover, as discussed in Section 3.1, the LB decision maker perceives no skewness in the distribution of the selling price of Security A. The true distribution of Y_T is skewed, however, and the LB model thus predicts a false assessment of the 10th and 90th quantiles. In particular, the LB decision maker fails to realize that the distribution's right tail is long, especially with a long investment horizon. We chose the experimental parameters such that under the LB model this bias would have no discernible effect in treatment HVS, but predicts an effect in HVL. The model predicts the 90-10 spread is too narrow-band in treatment HVL. (The LB model's lower bound point predictions are $q_{0.1,i}^{HVS} = 30$ and $q_{0.9,i}^{HVS} = 190$ in treatment HVS, and $q_{0.1,i}^{HVL} = 0$ and $q_{0.9,i}^{HVL} = 550$ in treatment HVL. Rational quantile assessments are identical in HVS but $q_{0.9}^{HVL}$ is at 700 in HVL.)

¹⁷All participants passed the understanding test, in a few cases after asking for some additional explanations.

¹⁸The instructions explain that violations of these constraints are subjectively suboptimal. Additionally, the experimental software shows an error message if a participant violates either of the two constraints. Only 2% of the participants' inputs receive one or more error messages.

¹⁹Participant can earn the bonus either through the quantile elicitation task or through the profit probability elicitation task (see Footnote 16). For each round and each participant, the relevant task type (quantile or profit probability) is determined by a simulated coin flip at the end of the experiment.

4.3 Results

We start our data analysis with a descriptive overview. Then, we run interval regressions that impose a normal decision error assumption and allow us to estimate the underlying quantiles while taking into account the data's discrete nature.

4.3.1 Descriptive Statistics

The boxplots in Figures 2 and 3 summarize participants' perceptions, separately for each treatment and for each of the five rounds. They depict the distributions of $q_{j,i}^k$ across all participants i , for the three probabilities $j \in \{0.1, 0.5, 0.9\}$ and for treatments $k \in \{\text{HVS}, \text{HVL}\}$. The horizontal dashed lines illustrate the benchmark rational predictions for the respective treatment-specific quantiles of the price distribution, q_j^k .

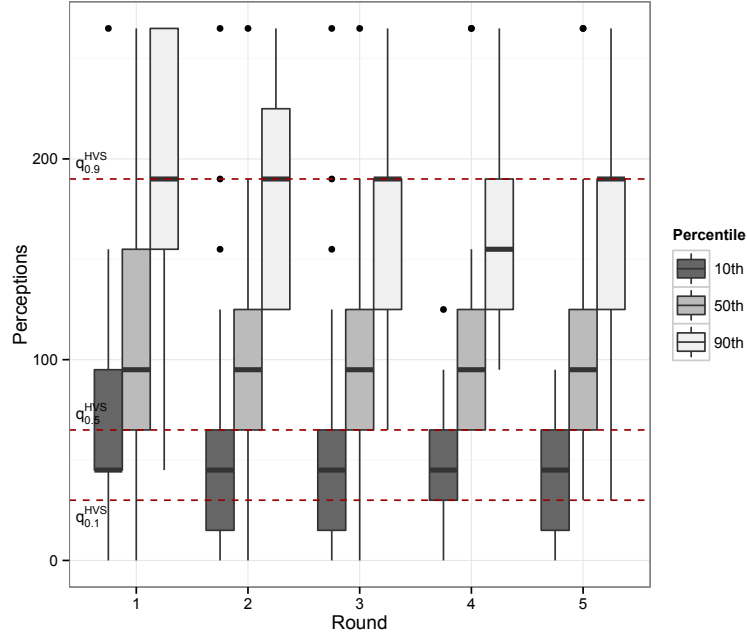


Figure 2: Elicited distributions of the participants' subjective quantiles of Security A's selling price in HVS, separately for the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed lines indicate rational benchmarks.

Figure 2 describes the switching points in treatment HVS (14-period time horizon). It shows that in all five rounds the median observation of $q_{0.5,i}^{HVS}$ (the solid line in the middle boxplot within each triplet of boxplots) is strictly above the rational benchmark $q_{0.5}^{HVS}$. Thus in each round at least half of the participants strictly overestimate the median of the stochastic process in treatment HVS. Precisely, the proportions of participants revealing that their subjective median is strictly too high are {Round 1: 55%, Round 2: 65%, Round 3: 55%, Round 4: 69%, Round 5: 62%}. Of all five rounds,

however, only Round 4's proportion is significantly greater than 50% (p-values<0.05, one-sided binomial test), which confirms our LB model prediction of an only marginal median misperception in this treatment.

The 90-10 spread, as perceived by the participants, can be roughly assessed from the location of the two boxplots depicting the subjective 10th (first boxplot within each triplet) and 90th percentiles (third boxplot within each triplet). As predicted by the LB model, these spread perceptions deviate only slightly from the rational prediction. This pattern, too, appears in each of the five rounds. In specific, the round wise proportions of participants individually underestimating the spread are {Round 1: 65%, Round 2: 65%, Round 3: 72%, Round 4: 72%, Round 5: 65%}. Of the five rounds, however, only in Round 3 and Round 4 are these proportions significantly greater than 50% (p-values<0.05, one-sided binomial test). With regards to perceived skewness, the arrangement of boxplots within one round also follows the LB model by being only marginally more symmetric than the rational prediction. Once again, we observe that the corresponding pattern is robust over the rounds.

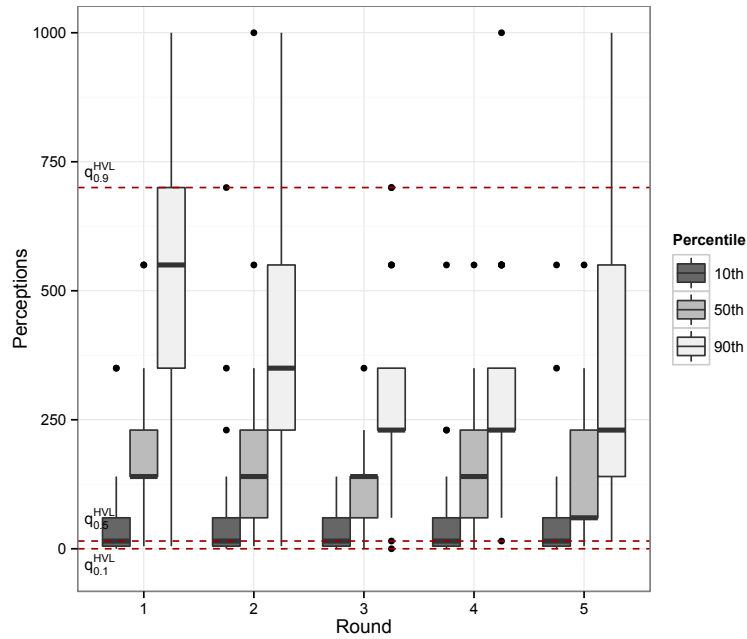


Figure 3: Elicited distributions of the participants' subjective quantiles of Security A's selling price in HVL, separately for the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed lines indicate rational benchmarks.

Supporting the LB model, we get qualitatively the same results for treatment HVL with the 140-period time horizon (Figure 3) but they are much stronger. Here, the interquartile ranges of the median-perception boxplots are located strictly above the optimal level in all rounds. Hence, more than 75% of participants overestimate the median in each round of this treatment. The precise round wise numbers are {Round 1: 93%, Round 2: 79%, Round 3: 79%, Round 4: 79%, Round 5: 79%}, all signif-

icantly greater than 50% (p-values<0.001, one-sided binomial test). Moreover, for all but two rounds are these proportions significantly greater in HVL than in HVS (p-value<0.05, two-sample binomial test). Likewise, the perceived price spread in this condition is underestimated by almost all participants over the rounds. In specific, the proportions for each round are {Round 1: 89%, Round 2: 96%, Round 3: 100%, Round 4: 96%, Round 5: 93%} (p-values<0.001, one-sided binomial test for >50%).²⁰ For each round these proportions are significantly greater in HVL than in HVS (p-value<0.05, two-sample binomial test). Moreover, the arrangement of boxplots within each round is obviously much more symmetric than the rational benchmark.

If we relax our rational benchmark a bit and count a participant as "overestimating the median" only if she switches more than one step (*task*) later than rational, we still count proportions of participants overestimating the median of 79% in the first round and 45% in the last round of treatment HVL. In HVS, these proportions are 34% and 37% for the first and last round, respectively.

4.3.2 Interval Regressions

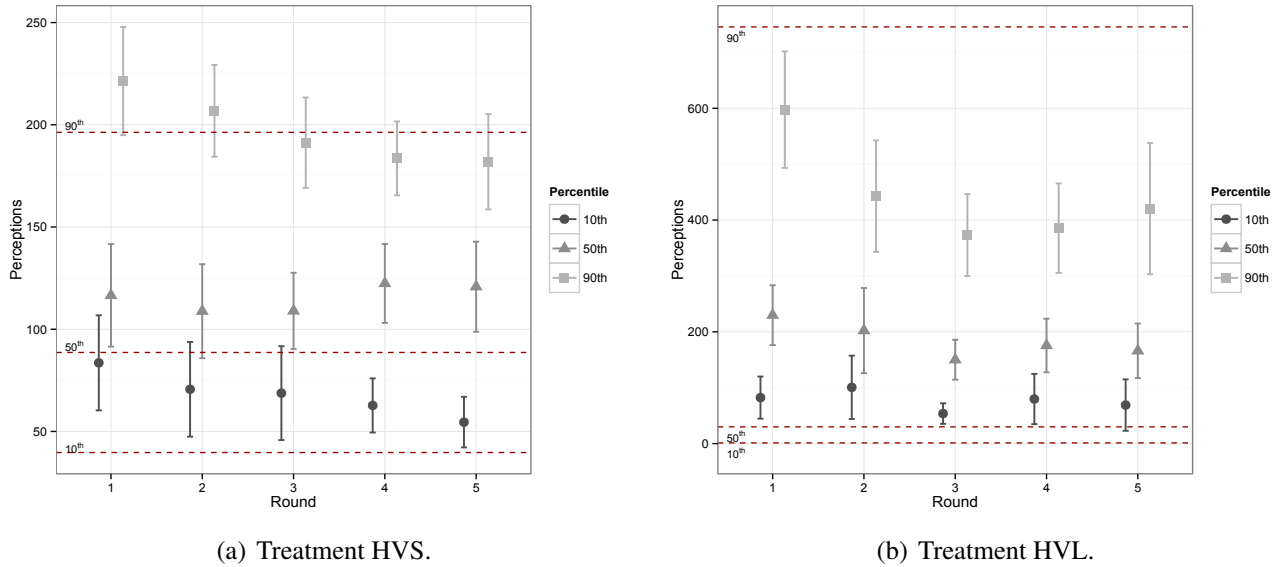


Figure 4: Point estimates enclosed by 95% confidence intervals of the participants' subjective quantiles of Security A's selling price. For each of the five rounds, separate estimates refer to the subjective 10th percentiles (circle), subjective medians (triangle), and subjective 90th percentiles (square). Dashed lines indicate rational benchmarks.

The above descriptive analysis is based on the elicited lower bounds of our participants' perceived quantiles of Security A's price distribution. This subsection investigates point estimates instead of

²⁰LB predictions and rational benchmarks for the 10th and 90th percentiles differ only for the latter in HVL. There, the proportions of those underestimating the 90th percentile range between 72% and 96% over the rounds and are always significant supporting the LB model (p-values<0.05, one-sided binomial test for >50%). In HVS, where LB and rational predictions coincide for the 90th percentiles, the respective proportions are much lower and not significant ranging between 31% and 58%.

lower bounds and employs interval regressions—a modified version of the ordered probit regressions (see e.g. Wooldridge, 2002). The analysis takes into account the interval nature of the data and assumes that the subjectively perceived quantiles are subject to normally distributed disturbances. Under this assumption, the mean of the participants’ subjective quantiles can be estimated via maximum likelihood, and standard hypothesis testing applies. Figure 4 reports the corresponding estimates of the population means for the subjectively perceived quantiles about Security A’s selling price, separately for each of the two treatments and for each round. The horizontal dashed lines depict the benchmark rational predictions for the respective treatment-specific quantiles. The point estimates of participants’ average perceptions are enclosed by the respective 95% confidence intervals.²¹

Figure 4 confirms the previous subsection’s descriptive analysis. In HVS (Figure 4(a)), the 95% confidence intervals of the point estimates for the subjective median are only marginally above the rational predictions. In accordance with the LB model, in HVL (Figure 4(b)) median perception confidence intervals differ significantly from the rational benchmark for all five rounds representing a much stronger degree of overestimation compared to HVS.

Overall, Figure 4 nicely sums up the results from Study 2, the elicited quantiles indicate that subjective skewness and spread perceptions of the outcome distribution tend to be quite accurate in HVS with a marginal tendency towards underestimation and this effect exacerbates considerably for a longer investment horizon represented by HVL.

5 Study 3: Linearity Bias in the Perception of Exchange Traded Funds

In this study we test the robustness of the LB model in a setting where the asset price depends on real-world historical data. We simulate leveraged and unleveraged exchange-traded funds (ETFs) on past data of the German stock market index DAX30 to examine how changes in volatility affect the participants’ perceptions of real-life growth processes.

Leveraged ETFs move by a given multiple relative to an underlying asset, compounded at the end of each trading day. A triple leveraged ETF on the DAX30 index increases by three per cent on a trading day if the DAX30 increases by one per cent on that day and it falls by three per cent if the DAX30 falls by one per cent. Leveraged ETFs are a popular asset class amongst household investors but have come under severe scrutiny as many investors were perplexed when the products made a loss in a period where the underlying index made a gain.²² Our experiment confirms the prediction derived from the LB model that skewness and spread are strongly underestimated if the volatility is high.

²¹For a detailed listing of interval regression estimates, see Appendix C.1.

²²Regulatory units and the financial media issued extensive warnings that involve explanations of these counter-intuitive possibilities. The U.S. securities regulator FINRA issued a note in 2009 (FINRA Regulatory Note 09-31) saying that "...while such products may be useful in some sophisticated trading strategies, they are highly complex financial instruments that are typically designed to achieve their stated objectives on a daily basis. Due to the effects of compounding, their performance over longer periods of time can differ significantly from their stated daily objective..."

Section 5.1 describes the experimental design. Section 5.2 presents the LB model’s predictions and section 5.3 discusses the results.

5.1 Experimental Design

There are two treatments in this study. In both treatments, Security A is an ETF based on the DAX30. The two treatments differ only in that their respective versions of Security A differ in per-period volatility. In treatment ETF_3, the relevant security is a triple-leveraged ETF based on the DAX30. Its price changes, on each trading day, by three times the daily percentage changes of the underlying index DAX30. In treatment ETF_1, in contrast, Security A is simply the DAX30 ETF itself.

The time horizon until maturity of the ETF is 2,000 trading days both for ETF_1 and ETF_3. To generate realized price paths for the two assets, we sample 2,000 consecutive DAX30 closing values, drawn at random from the time period 1964 to 2012.²³ Participants can buy the ETF at a price of 100; if they buy it, they have to hold it for 2,000 trading days. As in Study 2, in both treatments, a participant who chooses Security A receives a fixed bonus if the selling price exceeds a given threshold t_A . These thresholds do not differ between ETF_3 and ETF_1. They are listed in Table 4.

To elicit three different quantiles, the alternative choice option Security B has three different specifications which are equal to those in Study 2, i.e., Security B1 yields the bonus with 90% probability, B2 with 50% and B3 with 10%. Like in Study 2, each participant faces three choice lists and the choices allow us to infer the subjectively perceived quantiles.

Thresholds for Security A in ETF_3 and ETF_1	
<i>Task 1</i>	30
<i>Task 2</i>	60
<i>Task 3</i>	90
<i>Task 4</i>	140
<i>Task 5</i>	200
<i>Task 6</i>	260
<i>Task 7</i>	330
<i>Task 8</i>	450
<i>Task 9</i>	650
<i>Task 10</i>	1,000
<i>Task 11</i>	1,600

Table 4: The 11 thresholds.

The computer terminals report feedback to the participants in the form of a sample selling price of Security A. That is, the computer randomly samples a sequence of 2,000 consecutive trading days

²³The instruction in the ETF treatments are analogous to the other treatments in Study 2. Additionally, participants receive general information about the DAX30 and a data summary of daily DAX30 movements in the relevant time period. The information is given in the form of a histogram as well as statements specifying the 90% confidence interval $([-1.8\%, 1.8\%])$ and the overall average of daily percentage changes (0.03%). Note that the participants are UK-based students who typically have little knowledge about German stock markets.

from the set of all available 2,000-day histories of the DAX30 and uses it to simulate the asset price at maturity. The participants learn the result of the simulation and their payoff. All other aspects of the protocol are identical to Study 2. In both treatments of Study 3, the basic procedure is repeated four times, making for five identical rounds for each participant. 59 participants are in one of the treatments of Study 3 (29 in ETF_3 and 30 in ETF_1), all of them undergraduate students at University College London. The incentivisation structure contains a participation fee and a possible bonus of £ 5.00 for each round.

5.2 Linearity Bias Prediction

While Study 2 analysed the effects of an increased time horizon, the treatment comparison in Study 3 focuses on the effects of increasing the per-period volatility. We generate the perceived distribution of an LB decision maker by means of simulations: we randomly sample 500 price paths of ETF_3 and ETF_1 as perceived by an LB decision maker to generate the perceived distribution. By analogy to the previous discussion, the LB decision maker is assumed to correctly perceive the distribution of daily changes in the relevant asset price, but views them as an absolute changes and perceives their distribution to be constant over time. She therefore neglects all compounding. In detail, we simulate the perceived selling price of an ETF by first randomly selecting a start date $t = 0$ at which the price is fixed (by design of the experiment) at $Y_0 = 100$. For each of the ensuing 2,000 trading days $s > t$ we consider the relative change in value on that day, μ_s , and add $\tilde{\eta}_s = (\mu_s - 1)Y_0$ to the current perceived price of the asset: $\tilde{Y}_s = \tilde{Y}_{s-1} + \tilde{\eta}_s$ (with correct perception of the starting value, $\tilde{Y}_0 = Y_0$). For example, suppose that $t = \text{Jan 2, 1978}$ was randomly chosen as the starting date and that the DAX30 increased in value by 1.4% on $s = \text{Feb 15, 1979}$. As the perceived absolute increase on the latter date, the simulation simply adds $\tilde{\eta}_s = 1.4$ to the price of the asset, as perceived by the LB decision maker. This way the simulation arrives at a perceived selling price at maturity. Repeating this procedure 500 times for random starting dates generates the perceived distribution of selling prices.

Consistent with the results of previous sections, the simulations generate the predictions that the increased per-period volatility of leveraged ETFs leads to a larger bias concerning median overestimation and concerning the underestimation of the outcome distribution's skew and spread. This is intuitive as the neglect of compounding is worse in situations where the proportional growth rates are further away from 1 and their multiplication is thus more different from a perceived linear growth.

5.3 Results

Figure 5 illustrates the distributions of participants' perceptions in ETF_1. It shows that all medians of the median-perception distributions (the solid line in the middle boxplot within each triplet of boxplots) are above the optimal level. Participants in our experiment are overly optimistic about the simple index ETF. The arrangement of boxplots also shows that the perceived distributions are quite symmetric. But we note that with a simple ETF, the true distribution is relatively symmetric as well.

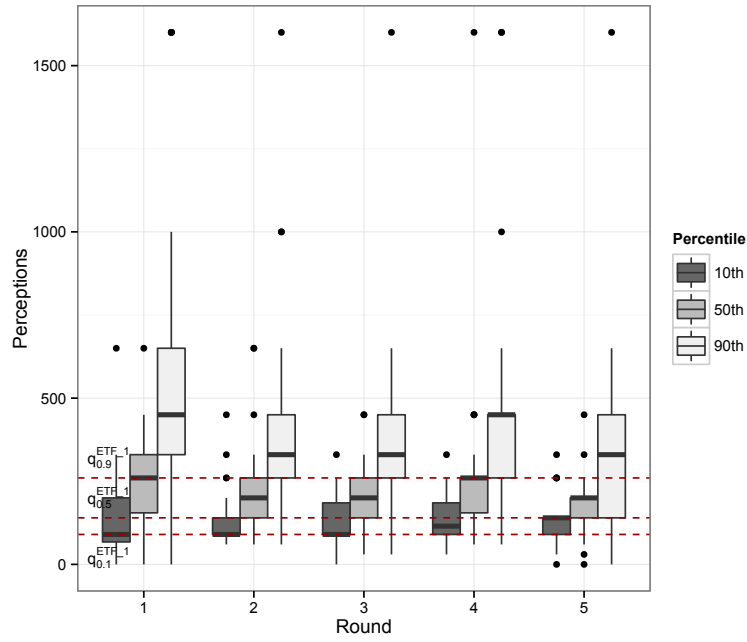


Figure 5: Elicited distributions of the participants' subjective quantiles of Security A's selling price in ETF_1, separately for the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed lines indicate rational benchmarks.

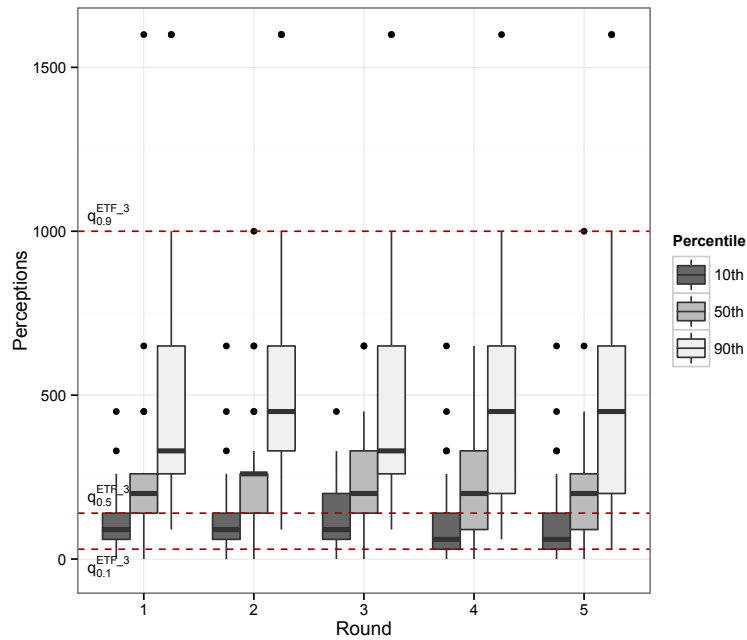


Figure 6: Elicited distributions of the participants' subjective quantiles of Security A's selling price in ETF_3, separately for the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed lines indicate rational benchmarks.

Figure 6 captures the participants’ perceptions in treatment ETF_3. Again, perceived medians show a notable level of overestimation. That is, the median of the perceived median distributions lies strictly above the optimal level. The arrangement of boxplots within a round also shows that the perceived spread and the perceived skewness are too small. The participants show at least a mild tendency to report skewed distributions but they far underappreciate the actual level of skewness.

We now consider interval regressions, in Figure 7. As before, they rely on the assumption of normally distributed decision errors, and they also confirm the main conclusions drawn from Figure 5 and 6. That is, in both treatments, ETF_1 and ETF_3, the outcome distributions’ median is significantly overestimated within each of the five rounds. Supporting the LB model, in treatment ETF_3, which has a much more skewed and spread out outcome distribution than ETF_1, participants far underappreciate the skewness and spread over all rounds.

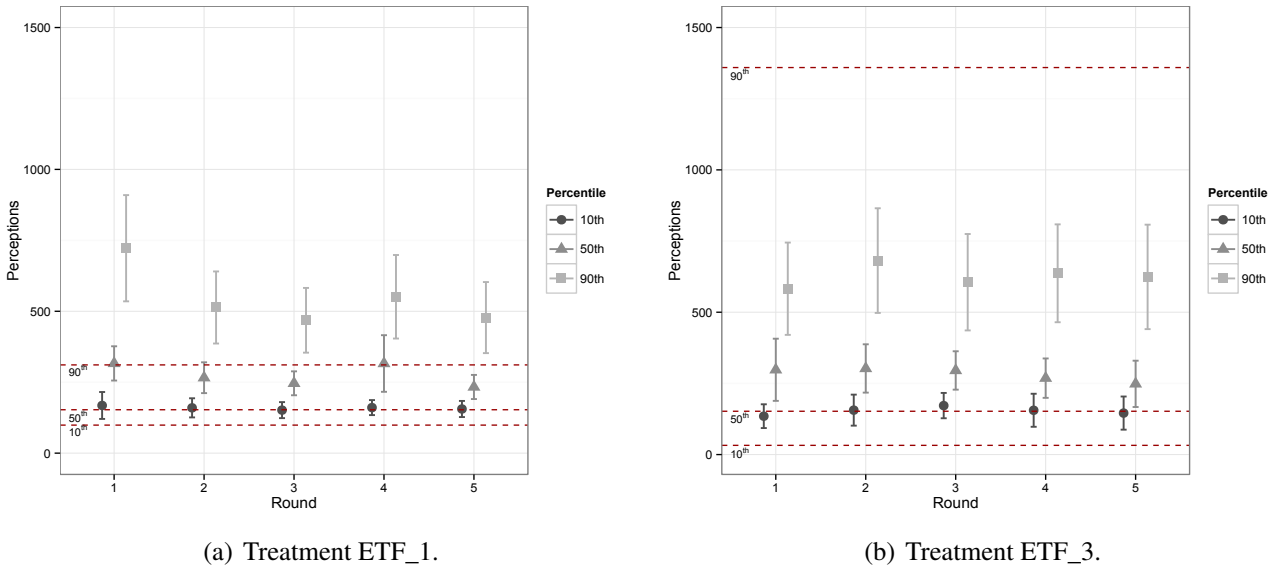


Figure 7: Point estimates enclosed by 95% confidence intervals of the participants’ subjective quantiles of Security A’s selling price. For each of the five rounds, separate estimates refer to the subjective 10th percentiles (circle), subjective medians (triangle), and subjective 90th percentiles (square). Dashed lines indicate rational benchmarks.

6 Conclusion

This paper investigates how people perceive two important implications of compounding random growth. First, as has been established in the literature, decision makers have a tendency to neglect exponential growth. Our experiments add to the evidence of this effect by measuring it in a context with random growth rates. Second, people underestimate the level of asymmetry in growth processes—skewness is “hidden”. This is a novel effect in the academic literature (modulo the independent description in Stutzer and Jung Grant (2010)), which may be especially relevant in the context of

preferences that consider quantiles of the outcome distribution, like value-at-risk or expected utility with highly convex utility functions. However, it is important to note that this paper is about misperception, not preferences, and that we measure the effects irrespective of risk attitudes.

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. Our experiments arguably give the respondents a very good shot at correctly detecting the nonlinear effects of growth, especially since we use highly selected and quantitatively skilled students and in some of our treatments we provide them with calculators. It is perhaps all the more notable that we, too, find a strong bias towards linear perceptions.

References

- [1] Abbink, K., and Rockenbach, B. (2006). *Option pricing by students and professional traders: a behavioral investigation*. *Managerial and Decision Economics* 27 (6): 497-510.
- [2] Benartzi, S., and Thaler, R.H. (1999). *Risk aversion or myopia? Choices in repeated gambles and retirement investments*. *Management Science* 45: 364-381.
- [3] Br  nner, T., Levinsky, R., and Qiu, J. (2011). *Preference for skewness: evidence from a binary choice experiment* *European Journal of Finance*, 17: 525-538.
- [4] Chen, H., and Rao, A.R. (2007). *When two plus two is not equal to four: Errors in processing multiple percentage changes*. *Journal of Consumer Research* 34: 327-340.
- [5] Christandl, F., and Fetchenhauer, D. (2009). *How laypeople and experts misperceive the effect of economic growth*. *Journal of Economic Psychology* 30: 381-92.
- [6] Cox, J.C., Ross, S.A., and Rubinstein, M. (1979). *Option pricing: a simplified approach*. *Journal of Financial Economics* 7: 229-263.
- [7] Deck, C., and Schlesinger, H. (2010). *Exploring higher order risk effects* *Review of Economic Studies* 77: 1403-1420.
- [8] Ebert, S., and Wiesen, D. (2011). *Testing for prudence and skewness-seeking* *Management Science* 57: 1334-1349.
- [9] Eckel, C.C., and Grossman, P.J. (2014) *Loving the long shot: risk taking with skewed lotteries*, Working paper, Texas A&M University.
- [10] Eisenstein, E.M., and Hoch, S.J. (2005). *Intuitive compounding: Framing, temporal perspective, and expertise*. Working paper, Johnson Graduate School of Management, Cornell University.
- [11] Ensthaler, L., Nottmeyer, O., Weizs  cker, G., and Zankiewicz, C. (2013) *Hidden skewness: on the difficulty of multiplicative compounding under random shocks*, Working paper, Humboldt University Berlin.
- [12] Financial Industry Regulatory Authority (2009), FINRA Regulatory Note 09-31:1-5
- [13] Fischbacher, U. (2007). *z-Tree: Zurich Toolbox for Ready-made Economic Experiments*, *Experimental Economics* 10(2): 171-178.
- [14] Gneezy, U. (1996). *Probability judgments in multi-stage problems: experimental evidence of systematic biases*. *Acta Psychologica* 93: 59-68.
- [15] Grether, D. (1981). *Financial Incentive Effects and Individual Decision Making* Working Paper 401, California Institute of Technology.

- [16] Holt, C.A. (2007). *Markets, Games & Strategic Behavior*.
- [17] Kemp, S. (1984). *Perception of changes in the cost of living*. Journal of Economic Psychology 5(4): 313-323.
- [18] Klos, A., Weber, E.U., and Weber, M. (2005). *Investment decisions and time horizon: risk perception and risk behavior in repeated gambles*. Management Science 51: 1777-1790.
- [19] Levy, M., and Tasoff, J. (2014). *Exponential-growth bias*. Working paper, London School of Economics and Political Science.
- [20] Lusardi, A., and Mitchell, O.S. (2011). *Financial literacy and planning: Implications for retirement wellbeing*. NBER Working Paper 17078.
- [21] McKenzie, C.R.M., and Liersch, M.J. (2011). *Misunderstanding savings growth: Implications for retirement savings behavior*. Journal of Marketing Research 48: S1-S13.
- [22] Redelmeier, D.A., and Tversky, A. (1992). *On the framing of multiple prospects*. Psychological Science 3(3): 191-93.
- [23] Samuelson, P. A. (1963). *Risk and uncertainty: a fallacy of large numbers*. Scientia 98: 108-113.
- [24] Stango, V., and Zinman, J. (2009). *Exponential growth bias and household finance*. Journal of Finance 64(6): 2807-2849.
- [25] Stutzer, M., and Jung Grant, S. (2010). *Expected return or growth rate? Choices in repeated gambles that model investments*. Working paper, Leeds School of Business, University of Colorado.
- [26] Wagenaar, W.A., and Sagaria, S.D. (1975). *Misperception of exponential growth*, Perception & Psychophysics 18(6): 416-422.
- [27] Wagenaar, W.A., and Timmers, H. (1978). *Extrapolation of exponential time series is not enhanced by having more data points*, Perception & Psychophysics 24(2): 182-184.
- [28] Wagenaar, W.A., and Timmers, H. (1979). *The pond-and-duckweed problem: Three experiments on the misperception of exponential growth*, Acta Psychologica 43(3): 239-251.
- [29] Wooldridge, J.M. (2002). *Econometric analysis of cross section and panel data*, Cambridge, MA: MIT Press.

A Study 1(c), Treatment Variation HELP

A.1 Treatment Description and Linearity Bias Prediction

Participants in Study 1(c) are randomly assigned to one of two treatments. Treatment NO_HELP (N=68), which is described in the main text above, presents only the basic explanation. In treatment HELP (N=60) 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 . None of the explanations in HELP adds any substantive information relative to the descriptions in NO_HELP. The only difference is that the relevant distributions are explicit in HELP and implicit in NO_HELP. Any difference in responses under the two conditions must stem from differences in the understanding of these implied truths.

Thus, in treatment HELP the linearity bias cannot influence the subjective beliefs without contradicting the available explanations. We therefore expect the misperception to disappear, i.e. $q_{0.5,i}^{HELP} = 500$.

A.2 Results

Share of participants switching from A to B					
Range of subjective median for Security A	Round 1	Round 2	Round 3	Round 4	Round 5
[0 – 100)	0.000	0.032	0.047	0.046	0.092
[100 – 500)	0.000	0.016	0.000	0.046	0.046
[500 – 2,000)	0.703	0.612	0.666	0.676	0.661
[2,000 – 6,000)	0.109	0.145	0.095	0.138	0.046
[6,000 – 9,000)	0.063	0.048	0.063	0.046	0.061
[9,000 – 12,000)	0.063	0.064	0.063	0.000	0.030
[12,000 – 20,000)	0.031	0.064	0.031	0.462	0.046
[20,000 – 35,000)	0.016	0.000	0.015	0.000	0.015
[35,000 – 90,000)	0.000	0.000	0.015	0.000	0.000
[90,000 – 250,000)	0.000	0.000	0.000	0.000	0.000
[250,000 – ∞)	0.016	0.016	0.000	0.000	0.000

Table 5: Subjective medians in HELP for rounds 1-5.

Table 5 shows that in HELP 70% of responses are at the optimal switching point of Task 3 already in round 1. Parametric t-tests as well as non-parametric Wilcoxon rank-sum tests confirm that all round-by-round treatment comparisons between HELP and NO_HELP are statistically significant at $p < 0.001$. Performance is poor under the NO_HELP condition and much better in HELP.

B Study 2, Low Volatility Treatment Variation

	Values of t_A <i>LVS</i>	Values of t_A <i>LVL</i>	Values of t_A <i>LVL_NC</i>
<i>Task 1</i>	104.0	185	185
<i>Task 2</i>	104.5	210	210
<i>Task 3</i>	105.5	240	240
<i>Task 4</i>	107.0	290	290
<i>Task 5</i>	109.0	340	340
<i>Task 6</i>	111.5	400	400
<i>Task 7</i>	114.5	460	460
<i>Task 8</i>	118.0	520	520
<i>Task 9</i>	122.0	625	625
<i>Task 10</i>	126.5	850	850

Table 6: The thresholds t_A by treatment condition.

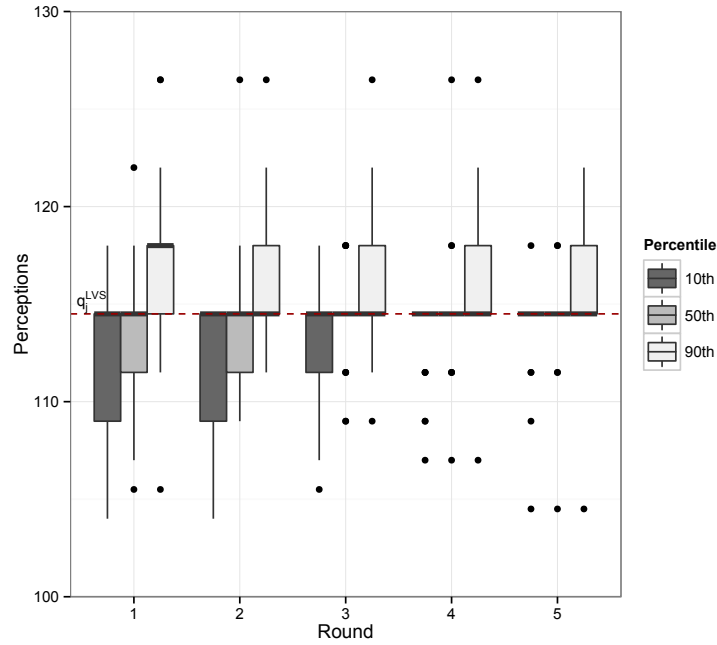


Figure 8: Participants' subjective quantiles of Security A's selling price in LVS over the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed line indicates rational benchmarks.

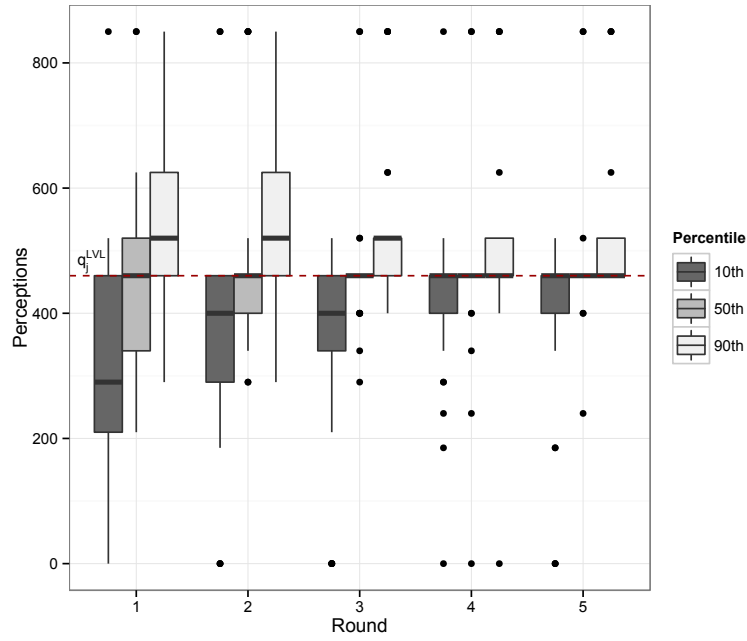


Figure 9: Participants' subjective quantiles of Security A's selling price in LVL over the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed line indicates rational benchmarks.

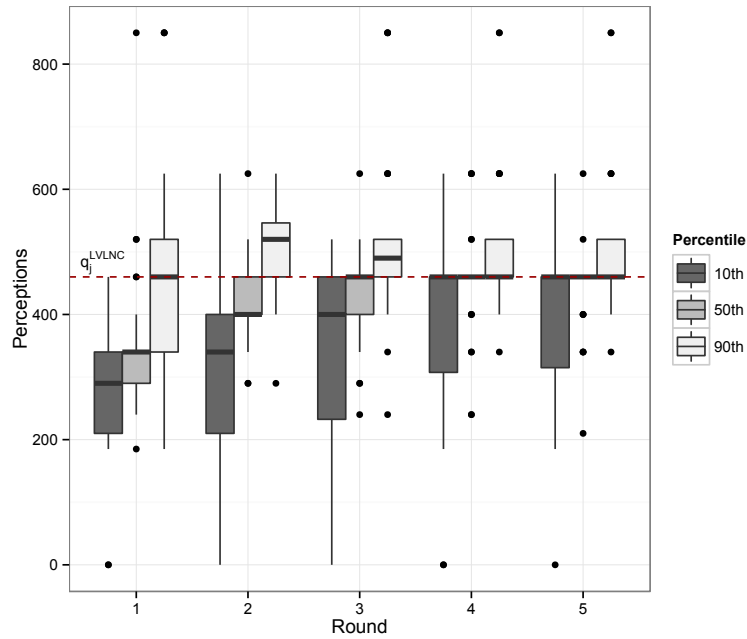
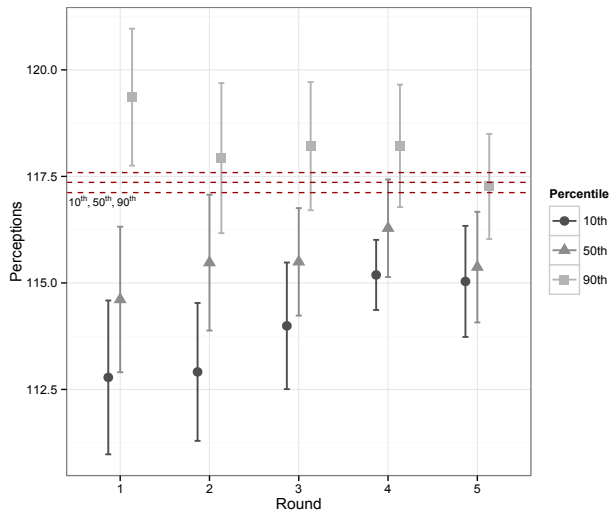
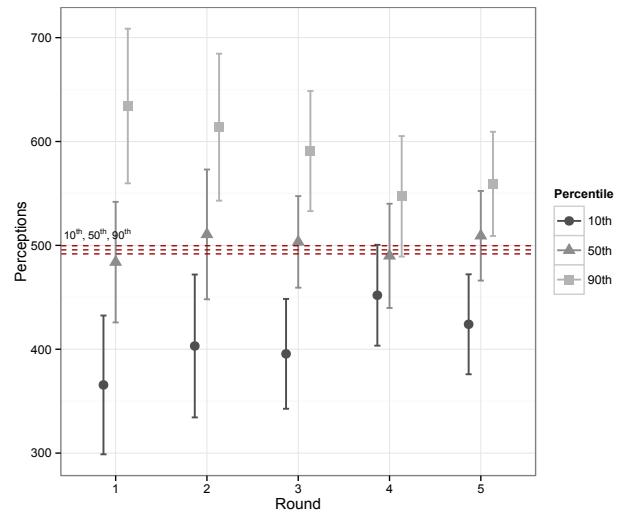


Figure 10: Participants' subjective quantiles of Security A's selling price in LVLNC over the five rounds. Each boxplot triplet consists of (from left to right): subjective 10th percentiles (dark grey boxplot), subjective medians (medium grey boxplot), and subjective 90th percentiles (light grey boxplot). Dashed line indicates rational benchmarks.



(a) Treatment LVS.



(b) Treatment LVL.

Figure 11: Point estimates enclosed by 95% confidence intervals of the participants' subjective quantiles over the five rounds, with subjective 10th percentiles (circle), subjective medians (triangle), and subjective 90th percentiles (square). Dashed lines indicate rational benchmarks.

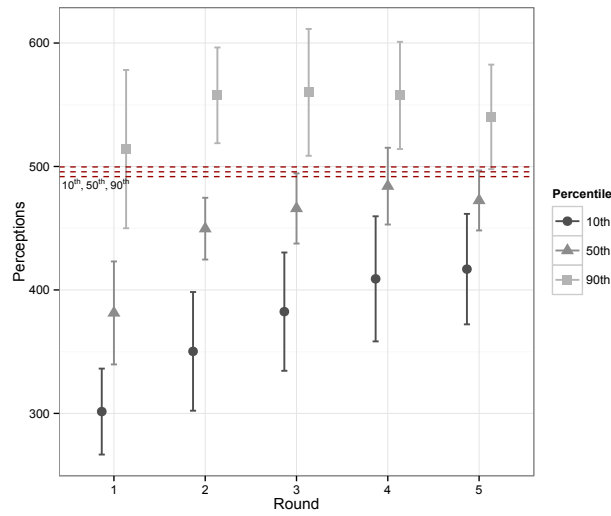


Figure 12: Point estimates in LVLNC enclosed by 95% confidence intervals of the participants' subjective quantiles over the five rounds, with subjective 10th percentiles (circle), subjective medians (triangle), and subjective 90th percentiles (square). Dashed lines indicate rational benchmarks.

C Tables

C.1 Interval regressions

	10th Percentile		50th Percentile		90th Percentile	
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Rational Model</i>	39.69		88.64		197.98	
<i>LB Model</i>	35.90		116.10		196.30	
<i>Round 1</i>	83.56	(11.8)	116.56	(12.7)	221.40	(13.5)
<i>Round 2</i>	70.63	(11.8)	108.84	(11.7)	206.84	(11.4)
<i>Round 3</i>	68.74	(11.7)	108.97	(9.5)	191.24	(11.2)
<i>Round 4</i>	62.72	(6.7)	122.41	(9.8)	183.57	(9.2)
<i>Round 5</i>	54.55	(6.3)	120.79	(11.2)	181.92	(11.9)

Table 7: Interval regression estimates for the mean perceptions of the three elicited percentiles in HVS complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Rational Model</i>	1.20		29.96		745.58	
<i>LB Model</i>	-59.80		261.00		581.80	
<i>Round 1</i>	82.27	(19.2)	229.80	(27.3)	597.49	(53.2)
<i>Round 2</i>	100.65	(28.9)	202.24	(38.9)	442.75	(50.8)
<i>Round 3</i>	53.68	(9.3)	150.07	(18.2)	373.06	(37.4)
<i>Round 4</i>	79.75	(22.9)	175.52	(24.4)	385.44	(40.9)
<i>Round 5</i>	68.90	(23.4)	165.97	(24.8)	420.60	(59.9)

Table 8: Interval regression estimates for the mean perceptions of the three elicited percentiles in HVL complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Rational Model</i>	117.12		117.36		117.59	
<i>LB Model</i>	115.90		116.10		116.30	
<i>Round 1</i>	112.78	(0.92)	114.61	(0.87)	119.36	(0.82)
<i>Round 2</i>	112.91	(0.83)	115.47	(0.81)	117.93	(0.90)
<i>Round 3</i>	113.99	(0.76)	115.49	(0.64)	118.21	(0.77)
<i>Round 4</i>	115.18	(0.42)	116.28	(0.58)	118.21	(0.73)
<i>Round 5</i>	115.03	(0.67)	115.37	(0.66)	117.26	(0.63)

Table 9: Interval regression estimates for the mean perceptions of the three elicited percentiles in LVS complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Rational Model</i>	491.79		495.69		499.63	
<i>LB Model</i>	260.20		261.00		261.80	
<i>Round 1</i>	365.65	(34.9)	483.87	(29.6)	634.18	(37.9)
<i>Round 2</i>	403.14	(35.1)	510.56	(31.8)	613.80	(36.1)
<i>Round 3</i>	395.55	(26.9)	503.36	(22.4)	590.81	(29.4)
<i>Round 4</i>	451.99	(24.7)	489.91	(25.6)	547.23	(29.5)
<i>Round 5</i>	424.01	(24.5)	509.23	(22.0)	559.20	(25.6)

Table 10: Interval regression estimates for the mean perceptions of the three elicited percentiles in LVL complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Rational Model</i>	491.79		495.69		499.63	
<i>LB Model</i>	260.20		261.00		261.80	
<i>Round 1</i>	301.48	(17.8)	381.39	(21.28)	514.08	(32.7)
<i>Round 2</i>	350.25	(24.5)	449.63	(12.76)	557.62	(19.7)
<i>Round 3</i>	382.42	(24.4)	466.01	(14.51)	560.03	(26.2)
<i>Round 4</i>	408.99	(25.8)	484.08	(15.86)	557.55	(22.1)
<i>Round 5</i>	416.87	(22.8)	472.47	(12.39)	540.17	(21.6)

Table 11: Interval regression estimates for the mean perceptions of the three elicited percentiles in LVLNC complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Sim.Value	Bootstr.Conf.Inter.	Sim.Value	Bootstr.Conf.Inter.	Sim.Value	Bootstr.Conf.Inter.
<i>Rational Model</i>	98.80	[94.95; 101.25]	152.98	[146.18; 162.40]	310.98	[292.75; 322.40]
<i>LB Model</i>	109.36	[104.21; 113.83]	162.72	[156.63; 169.16]	227.10	[221.59; 230.66]
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Round 1</i>	168.02	(24.2)	316.25	(30.8)	722.22	(95.5)
<i>Round 2</i>	159.65	(17.2)	265.84	(27.7)	513.36	(64.8)
<i>Round 3</i>	151.62	(14.4)	246.15	(21.5)	468.22	(58.2)
<i>Round 4</i>	160.56	(13.5)	315.93	(50.9)	551.12	(75.1)
<i>Round 5</i>	155.61	(14.5)	233.18	(21.7)	477.73	(63.9)

Table 12: Interval regression estimates for the mean perceptions of the three elicited percentiles in ETF_1 complemented by LB and rational predictions at the top of each column.

	10th Percentile		50th Percentile		90th Percentile	
	Sim.Value	Bootstr.Conf.Inter.	Sim.Value	Bootstr.Conf.Inter.	Sim.Value	Bootstr.Conf.Inter.
<i>Rational Model</i>	32.39	[29.24; 36.31]	152.14	[123.17; 178.88]	1,359.42	[1,163.44; 1,555.64]
<i>LB Model</i>	128.09	[112.64; 141.51]	288.17	[269.91; 307.49]	481.30	[464.79; 491.98]
	Mean	Std.Err.	Mean	Std.Err.	Mean	Std.Err.
<i>Round 1</i>	134.87	(21.3)	297.83	(55.7)	582.54	(82.7)
<i>Round 2</i>	156.07	(27.8)	302.48	(43.3)	681.32	(93.8)
<i>Round 3</i>	172.03	(22.7)	295.44	(34.4)	605.39	(86.4)
<i>Round 4</i>	155.59	(29.6)	268.41	(35.3)	636.71	(87.7)
<i>Round 5</i>	145.72	(29.7)	248.39	(41.5)	624.07	(93.6)

Table 13: Interval regression estimates for the mean perceptions of the three elicited percentiles in ETF_3 complemented by LB and rational predictions at the top of each column.