

Dynamic Regret Avoidance[†]

By MICHELE FIORETTI, ALEXANDER VOSTROKNUTOV, AND GIORGIO CORICELLI*

In a stock market experiment, we examine how regret avoidance influences the decision to sell an asset while its price changes over time. Participants know beforehand whether they will observe the future prices after they sell the asset or not. Without future prices, participants are affected only by regret about previously observed high prices (past regret), but when future prices are available, they also avoid regret about expected after-sale high prices (future regret). Moreover, as the relative sizes of past and future regret change, participants dynamically switch between them. This demonstrates how multiple reference points dynamically influence sales. (JEL C91, G12, G41)

Regret is a negative emotion, associated with an action or inaction, that is experienced when one wishes that another choice would have been made. Regret avoidance was found to be an important factor in many empirical studies on topics ranging from heart disease prevention in health economics (Boeri et al. 2013) to auctions (Filiz-Ozbay and Ozbay 2007; Hayashi and Yoshimoto 2016), financial markets (Fogel and Berry 2006; Frydman, Hartzmark, and Solomon 2018; Frydman and Camerer 2016), portfolio and pension scheme selection (Muermann, Mitchell, and Volkman 2006; Hazan and Kale 2015), and currency hedging (Michenaud and Solnik 2008).

Apart from the empirical applications, regret avoidance has been studied both theoretically (Bell 1982; Loomes and Sugden 1982; Skiadas 1997; Sarver 2008; Hayashi 2008; Bikhchandani and Segal 2014; Leung and Halpern 2015; Qin 2015; Buturak

* Fioretti: Department of Economics, Sciences Po (email: michele.fioretti@sciencespo.fr); Vostroknutov: Department of Economics (MPE), Maastricht University (email: a.vostroknutov@maastrichtuniversity.nl); Coricelli: Department of Economics, University of Southern California (email: giorgio.coricelli@usc.edu). Michael Ostrovsky was coeditor for this article. We thank two anonymous referees for helping us improve the paper substantially. We are very much indebted to Geert Ridder for his insights and support on this project. We would also like to thank Martina Puppi and Folco Panizza for invaluable help with running the experiments, and seminar and conference participants at Alicante, Caltech, CIMEC, GMU, 2017 AMES, 2017 ESA World Meeting, 2017 NAMS, M-BEES 2017, 2017 ESA European meeting, 11th IMPRS Uncertainty Topics Workshop, Trento. This paper was greatly improved through conversations with Marina Agranov, Nadège Bault, Daniel Benjamin, Fanny Camara, Khai Chiong, Cary Frydman, Matthew Kahn, Chad Kendall, Erik Kimbrough, Yilmaz Kocer, Scott Kominers, Tobias Larsen, Michael Leung, Jacopo Magnani, Denis Nekipelov, Luca Polonio, Alejandro Robinson Cortes, Joel Sobel, Jorge Tamayo, Michael Woodford, Leeat Yariv, and Anthony Lee Zhang. The study was funded by the European Research Council (ERC Consolidator Grant 617629). Michele Fioretti gratefully acknowledges USC Dornsife INET for funding and support during the first stage of this project. Computation for the work described in this paper was supported by the University of Southern California Center for High-Performance Computing (hpc.usc.edu). All mistakes are our own. The IRB approval for the data collection was provided by University of Trento (Department of Economics).

[†] Go to <https://doi.org/10.1257/mic.20180260> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

and Evren 2017) and experimentally (Coricelli et al. 2005; Camille et al. 2004; Zeelenberg 1999; Bleichrodt, Cillo, and Diecidue 2010; Strack and Viefers 2021). Even though many aspects of regret avoidance were considered in these studies, their focus is mainly on *static* problems where a single decision is made that can be affected by the information about possible counterfactual outcomes. Such problems are important since many real-life decisions, like buying a house or a pension plan, fit into this setting. Nevertheless, many interesting phenomena that involve regret have *dynamic* nature, the stock market being one important example. These situations are characterized by the presence of the time dimension: a decision or decisions should be made given some past information and/or expectations of the future, both of which change as time unfolds. Regret, in this case, also becomes a dynamic variable that is reevaluated in each time period. More importantly, there emerge the concepts of *past* and *future regret*. A choice is influenced by past regret when an action taken today increases the chances of bringing about a desirable outcome that was observed in the past. Future regret involves taking actions that prevent missing the opportunity of achieving a desirable expected future outcome. For example, in financial markets, the decision to sell an asset might depend on the highest observed price in the past (past regret), but traders might also think about the hypothetical counterfactual situation in which they sell an asset today and regret doing it later because the price went up (future regret) and adjust their behavior to avoid such circumstances.

In this paper, we investigate how past and future regret influence choices in a controlled experimental setting similar to a stock market. Our main interest is to understand how different elements of the dynamic situation interact and influence behavior: in our case, the decision to sell an asset. In particular, we are interested in the following questions: (i) How strongly does the avoidance of past and future regret influence the choice to sell? (ii) Is there an interaction between past and future regret? Does one become stronger or weaker in the presence of the other?

In our experiment, reminiscent of those reported in Oprea, Friedman, and Anderson (2009); Oprea (2014); and Strack and Viefers (2021), participants take part in a series of “stock markets”: they observe how the price changes in real time and choose when to sell an asset that they own. Participants make choices in two types of markets. In some markets, they do not see the future price of the asset after they made their selling decision. In other markets, they do see the future price. Participants are always informed beforehand about the type of the market they are in. This setup allows us to analyze past and future regret and their interaction. In both conditions, past regret can potentially influence participants’ decisions to sell the asset since the price history is observable. At the same time, we are able to see whether access to the prices after selling has an effect on decision-making (future regret). More importantly, our design makes it possible to use structural modeling and estimate the parameters of a utility specification that includes past and future regret components in a dynamic discrete choice setting (e.g., Rust 1987; Hotz and Miller 1993).

We find that participants *are* influenced by the observable past prices and *do* behave differently depending on whether they know that the future prices will or will not be observed after they sell the asset. Our evidence that participants keep the asset to make the effect of past regret smaller or absent confirms the results of

the recent studies that focus on past regret only (Gneezy 2005; Strack and Viefers 2021). We go further and consider the possibility that agents keep the asset longer when they know that they can observe future prices and expect them to be high, as compared to the case when they know they will not observe future prices. Our data show that information about the availability of the prices after selling indeed has this expected effect on the decision to sell. More importantly, when the participants know that they will not observe future prices, their choices to sell are *not* affected by future regret avoidance. In addition, we find that individual risk preferences also play a role in the selling decisions. However, their effect on choice is secondary to regret avoidance and does not influence the estimates of the regret parameters.

Estimates of the parameters of a regret-averse utility function obtained from a dynamic discrete choice model suggest that the effects of the past and future regret are not simply additive. We demonstrate that there is an interaction between past and future regret in the utility, which would not be possible to identify with simple regression analysis. Past and future regret are not complements but rather lessen the effects of one another. This happens because, while both regret components of the utility function are negative, the interaction term offsets the effect of the smaller one. We call this phenomenon a *substitution effect* between past and future regret. At each point in time, participants' selling choices are not influenced by both types of regret at once but are rather guided by the one that is stronger. This also implies that depending on the circumstances, the behavior on the market can be either past or future oriented.

Our findings demonstrate that individuals incorporate past and future regret into the utility function in dynamic settings and that they are able to extract and update complex counterfactual information about the changing environment and integrate it into the decision process.

I. The Experiment

The data were collected in a behavioral experiment in which participants were presented with a series of mini stock markets. Each participant observed the graph of a market price as it gradually changed in time in 0.8-second intervals and had to decide when to sell an "asset" (see Figure 1). For the first 15 periods, participants could only observe the price.¹ Then, in period 15, they were forced to buy an asset at the current price. The point of entry was marked with a vertical red line. The market price kept changing until participants decided to sell the asset (marked with a blue line on the graph). In case no selling decision was made, the market continued until its closure in period 50, at which point participants were forced to sell. The profit was equal to the selling price minus the entry price (price in period 15) so that participants could actually lose money (each participant received a €10 fee that covered her in the case of a loss).

¹We included 15 initial nonchoice periods following Frydman and Camerer (2016), who also included nonchoice periods in a similar design to allow participants to experience the price variation within a market before making any decisions. Not including these periods can lead to a situation where participants stay in the market simply because they want to learn more about the price and not because of regret avoidance.



FIGURE 1. SCREENSHOTS OF TWO MARKETS

Notes: Above the graph, participants could see the entry price (valore di entrata), current price (valore corrente), selling price (valore di uscita), and profit (guadagno), which was green for positive profit and red for negative profit. In the No Info condition, the future price was not shown (left picture). In the Info condition, the price evolution was shown after the selling decision (right picture). The sentence at the bottom of the left picture says “Please wait until the market is closed.”

In each market, the price followed a stochastic mean reverting process defined by $y_{t+1} = \alpha y_t + (1 - \alpha)\theta$, where $\alpha = 0.6$, y_t is the price in period t , and θ is an identically and independently distributed random variable (uniform between €0 and €10). Participants were informed about the process that generated the price and made selling decisions in six training markets without payment, which allowed them to see the examples of the price dynamics and get used to the interface (the market prices used in the experiment are graphed in online Appendix A.5).

Each participant made selling decisions, which could be of two types, in 48 different markets. In some markets (No Info condition, left picture in Figure 1), participants *knew* from the beginning that after they sell the asset, they will not see the future price. In the Info condition (right picture in Figure 1), participants knew from the beginning that after selling the asset, they will observe the evolution of the price until the market closure in period 50. This information was shown in the upper-left corner of the graph from period 1 onward (INFO DOPO means “info after”). The markets were presented in a random order that was generated independently for each participant. Half of the markets were presented in the No Info condition and half in the Info condition. The sequence of conditions was also randomized. After the markets, the participants were presented with an incentivized Holt-Laury task (Holt and Laury 2002) and a questionnaire. Overall, 154 participants took part in the experiment in 9 sessions. The average earnings in the main task were €11.46. The experiment was programmed in z-Tree (Fischbacher 2007). The data and the analysis can be found at the data repository openicpsr-130441 (<http://openicpsr.org/>). Further details of the design can be found in online Appendix A.

II. Evidence of Regret Avoidance

In this section, we look at some summary statistics in order to compare the selling behavior to the no-regret benchmark, and we report a regression analysis that shows the effects of past and future regret. This analysis can provide only crude estimates of how the current market state influences the choices to sell, since it is




FIGURE 2. OPTIMAL SELLING PRICE THRESHOLDS

Notes: The optimal selling price threshold for different risk preferences. CRRA utility without regret is shown.

static in nature and does not take into account the dynamic structure of the markets. Nevertheless, it does demonstrate how the participants react to past and expected future prices. We start with a comparison of the behavior of our participants with the optimal choice of a risk-neutral regret-free agent who should sell the asset whenever the price rises above a certain threshold that depends on the number of periods left in the market. The dynamic stopping problem that describes optimal choices is formulated in online Appendix B.

We focus on the class of CRRA utilities $U(y) = (y^{1-\rho} - 1)/(1 - \rho)$, where y is a selling price, and numerically evaluate the optimal policy prescribed by the dynamic program from online Appendix B. Figure 2 illustrates the optimal policies for five values of the risk parameter ρ (both risk loving and risk averse).² It is optimal for the agent to sell the asset if the price is above the shown thresholds. The figure demonstrates that risk-loving agents (with $\rho < 0$) optimally sell the asset at higher prices than risk-averse agents ($\rho > 0$). Notice, however, that the effect of risk preferences on the optimal threshold is rather small. The threshold is virtually the same in period 16 for risk-loving and risk-averse agents, and in period 49, the threshold changes from €4.7 ($\rho = 0.7$) to €5.3 ($\rho = -1$). This implies that we should not expect any strong behavioral effects to stem from the heterogeneity in risk preferences.

In order to compare the behavior of participants with this benchmark, we consider selling decisions at relatively high prices, since participants' choices coincide with the model prediction to keep the asset when the prices are low. Figure 3 summarizes selling decisions in situations when the participants had a choice to sell

²We consider the values of ρ in the interval $[-1, 0.7]$ following Strack and Viefers (2021), who found that the estimates of CRRA risk coefficients for their subjects lie in this interval.




FIGURE 3. THE PROPORTION OF TIMES THE PARTICIPANTS DECIDED TO KEEP THE ASSET

Notes: The participants keep the asset despite the current price being greater than the optimal selling threshold of the risk-neutral regret-free agent. Observations are grouped by deciles. The solid line at zero shows the proportion of missed sales expected from the rational no-regret agent. The spikes are ± 1 SE.

and the price was above the optimal selling threshold (of a risk-neutral agent). If our participants chose in accordance with the predictions of the no-regret utility model, they would have sold the asset in *all* these cases. We observe that even at the tenth decile of the price distribution, there is a large deviation from the optimal strategy: participants do not sell the asset in 34 percent of the cases. When we look at the prices below €6.3 (first decile), we see that the asset is kept in 80 percent of the cases when it actually should have been sold, a huge discrepancy with the predictions of the standard model. Still, the deviations from the standard model can, in principle, be noise artifacts. To falsify this idea, we run a logit and an OLS regression where the dependent variable is the decision to keep the asset and the independent variable is the current price (conditional on being above the optimal threshold). We find a significant negative trend in the probability to keep the asset (logit coefficient -0.71 , OLS coefficient -0.16). This shows that the differences in proportions are not random and are higher for lower prices. Finally, to get an idea about the between-subject differences in behavior, we plot on Figure C4 in online Appendix C the distribution of participants by their proportion of missed optimal sales calculated for each participant separately. We see that there are very few participants who are close to the proportion of zero predicted by the no-regret model. Only around 5 percent of participants keep the asset when it should be sold in less than 30 percent of cases; everyone else keeps the asset in more than 30 percent of cases. This demonstrates that very few participants, if any, are behaving in accordance with the no-regret model.

This evidence suggests that the participants mostly keep the asset in situations when the standard model predicts that it should be sold. One potential explanation of this effect is loss aversion. Suppose that participants suffer some additional fixed disutility from having negative profit (selling the asset at a price lower than the entry price). This can, in principle, make them keep the asset longer in order to make a positive profit. However, in our data, the correlation between the entry and selling prices is very small (Spearman's $\rho = 0.058$, $p < 0.001$). Moreover, in Figure 3, the average proportion of missed sales over all price deciles is 0.58. If we only look at the data points where participants would have made positive profit by selling, this proportion drops to 0.48, which is only 10 percent less. This means that when participants should sell the asset according to the no-regret model and can make positive profit, they still do not do it in 48 percent of the cases. All this evidence suggests that loss aversion is not a good candidate for explaining the data. Moreover, it does not predict any difference between the Info and No Info conditions, which we report below.

We hypothesize that the observed behavior is driven by the desire to minimize regret, which can arise in our dynamic setting due to two kinds of (counterfactual) comparisons between the outcome of a current choice (i.e., realized price) and past or future peaks. One possibility is that the decision-maker can take all observable past prices and form an expectation using this information about how high the price can go. In this case, she avoids what we call past regret by keeping the asset if the past information suggests that the price can increase further. We propose that the decision-maker focus on the highest price in the past, or past peak, to form this expectation.³ The past peak is calculated as the highest price achieved up to the current period.⁴ Another possibility is that the decision-maker anticipates regret from knowing that a higher price can be attained in case postsale prices are observable (the Info condition). So the decision-maker avoids future regret when she keeps the asset longer due to expectations that the price can increase after she sells it. Importantly, unlike the past prices that are always observable, future regret should only be relevant if the decision-maker knows that the future prices will be revealed, because there is no possibility to experience regret due to higher future prices otherwise (the No Info condition). Similarly to the past regret, we assume that the decision-maker uses the future expected highest price (future peak) as a reference point in this case. The future peak is computed as the expectation over the maximum price that can be achieved in all future paths (see Section III for details).

To test the idea that the past peak influences the decision to sell, we examine the selling rates. Overall, participants sell the asset in 51 percent of the situations when

³Gneezy (2005) shows in a setting similar to ours that the past peak is a more plausible reference point than the purchase price.

⁴Under the standard definition, regret is elicited by the counterfactual comparison between the outcome of the chosen and the outcome of a foregone option. This suggests that the definition of past regret should exclude the first 15 nonchoice periods in the experiment. However, in our specific setting related to the stock market, the decision-maker can form better price expectations when using all observable past prices regardless of whether they are before or after period 15. Thus, providing price information for the first 15 periods allows the decision-maker to anticipate regret based on past information starting already from her first choice in period 16. In our setting, extending the standard definition of regret to include the first 15 nonchoice periods is crucial to identify how past regret affects sales.




FIGURE 4. EVIDENCE OF REGRET AVOIDANCE

Notes: Panel A: The percentage of sales when the price reaches a new peak (dark gray) and when the price is below the current past peak (light gray). The error bars are ± 1 SE. Panel B: The ratio of the number of sales up to period t in the No Info condition to the number of sales up to period t in the Info condition. The dashed line includes sales within €0.5 of the past peak and the solid line within €1 of the past peak.

the price goes above the past peak (i.e., it is a new peak). If we look only at the cases when the price is above the optimal threshold of the risk-neutral agent without regret, then the selling rate becomes 71 percent at the new peak and 35 percent when the price is not the new peak. This provides evidence that the past peak has an influence on the decisions to sell even when the standard model unambiguously predicts only sales. To see the importance of the past peak for the decision to sell, consider Figure 4, panel A. We group the new past peaks by how high they are and find that when the new past peak is above €8, selling happens in 71 percent of the cases; in the range [7, 8], 63 percent; in the range [6, 7], 30 percent; and in the range [5, 6], 2.6 percent. Notice that the percentages of selling when the price is in the same intervals but is *not* a new peak are 46 percent, 32 percent, 14 percent, and 3.4 percent, respectively—much lower values. Figure C5 in online Appendix C shows the same graph restricted to observations above the optimal threshold in the standard model. The influence of the past peak is unchanged. Thus, the difference in sale rates cannot be explained by the standard theory; we need to consider past peaks in order to explain our data.

To show the influence of the future peak on the decisions to sell, we notice that the future expected highest price is decreasing in time, since early in the market, there are plenty of opportunities for the price to rise, whereas when there are only few periods left, the price cannot go much higher than its current level. Therefore, future regret, which is proportional to the future peak, should be highest in early periods and decrease later on. If our participants are sensitive to future regret, we should observe a difference in selling behavior between the No Info and Info conditions in early periods. The two curves on Figure 4, panel B represent the cumulative ratio of the number of sales in the two conditions that are within €0.5 and €1 of the past peak. For each time period, this ratio exceeds 1, which implies that there are more decisions to sell in the No Info condition than in the Info condition. This effect is especially evident in the early periods. In the late periods, the number of selling

TABLE 1—RANDOM EFFECTS LOGIT REGRESSION OF THE CHOICE TO KEEP THE ASSET

Pr[choice = keep]	I	II	III
Price	−0.497 (0.146)	−0.319 (0.133)	−0.326 (0.134)
Price ²	−0.102 (0.013)	−0.125 (0.012)	−0.125 (0.012)
Time	−0.088 (0.004)	−0.082 (0.004)	−0.082 (0.004)
Future expected price	1.423 (0.230)	1.401 (0.190)	1.381 (0.188)
Past peak		0.506 (0.035)	0.600 (0.045)
Future expected peak		0.309 (0.071)	0.210 (0.084)
Past peak × info			−0.209 (0.065)
Future expected peak × info			0.183 (0.066)
Info			0.129 (0.675)
Constant	4.525 (1.161)	−1.955 (0.966)	−1.746 (1.099)
Observations	112,137	112,137	112,137

Notes: Choice is 0 at the time the participant sells the asset and 1 otherwise. Observations are for all periods in all markets for all participants in which they made a choice (periods 16 to 49). Errors are clustered by participant and robust. The variable price refers to the current price; time is the period counter; future expected price refers to the expected price in period 50 given the current price, the number of periods left, and the price generating Markov process; past peak is the highest price observed before current period; future expected peak is the highest expected price given the current price and the number of periods left; and info is 1 if the condition is info and 0 otherwise. The descriptions of all variables can be found in online Appendix D.

decisions becomes approximately the same.⁵ This provides first evidence that participants sell less often early on in the Info condition because of the possibility of future regret, which makes them keep the asset longer in order to reduce the disutility associated with it. Moreover, the ratio is higher for the sales that are €0.5 closer to the past peak than for the sales that are €1 closer. This is the case since in the No Info condition, being closer to the past peak implies higher probability of selling, whereas in the Info condition, the past peak is less salient due to the possibility of observing high prices after selling.

To investigate the influence of a larger set of variables on the decisions to sell, we run a series of logit regressions, shown in Table 1, with the dependent variable equal to 1 if a participant keeps the asset and 0 if she sells it. Notice that these regressions can provide only a simplified picture of the relationships in our data since they do not account for the time dependencies due to the Markovian nature of the price

⁵Figure C6 in online Appendix C shows that the ratios starting from period 33 oscillate in the vicinity of 1.

evolution and the optimizing behavior of the participants. The main variables of interest are the market condition (info), the past peak, the future expected peak, and their interactions.⁶ We see that both past and future peaks significantly influence the probability to keep the asset (columns II and III): the higher they are, the longer the participants hold the asset. More importantly, in the Info condition, we see that the influence of the past peak decreases, and the influence of the future expected peak increases (interactions with the variable info, column III). This is consistent with our hypothesis that the possibility to observe prices after selling the asset makes participants more focused on the future. All these findings are in line with what the standard regret theories (e.g., Bell 1982; Loomes and Sugden 1982) would suggest in our setting. Namely, higher past and future expected peaks decrease the utility from selling the asset today, which results in the participants' holding on to it longer. This happens because these peaks suggest where the price can potentially be in the future if the asset is kept, thus creating regret associated with selling it today.⁷

We further investigate the decision to keep the asset by introducing more variables. The regressions reported in Table E2 in online Appendix E show a small but significant effect of the risk preferences, as estimated by the Holt and Laury task (Holt and Laury 2002), on the probability to keep the asset (variable hl). As risk aversion increases, the probability of keeping the asset goes down, which is consistent with the predictions of the no-regret model (online Appendix B). Nevertheless, risk preferences alone cannot account for the dependency of the selling choices on the market condition, the past price history, or future expected prices since all the interactions of the corresponding variables with hl are insignificant (regressions in columns V and VI). Finally, the regressions in Table E3 in online Appendix E show a significant effect of the market condition (Info versus No Info) in early periods. The probability of keeping the asset is higher in the Info condition (variable info \times early, columns I and II), which is in line with future regret avoidance, as we explained above (Figure 4, panel B).

To summarize, we find some evidence that the decisions to sell are influenced by past regret avoidance (Figure 4, panel A and Table 1). We also find that in the Info condition, future peaks become more and past peaks less salient for the decision to keep the asset (Table 1), which is consistent with future regret avoidance. Finally, participants keep the asset longer in the early periods of the Info condition when future regret is the strongest (Figure 4, panel A and Table E3 in online Appendix E), suggesting an interaction between the two types of regret. These results, however, should be treated with caution. While the presence of past regret avoidance is unambiguous, given that past peaks are always observed by participants and are in a sense "tangible," we cannot reliably conclude from the regression

⁶See online Appendix D for the description of the variables used in the regressions and online Appendix F for the computation of highest expected future price.

⁷We also make several observations about the control variables. The probability of selling increases with time (coefficient on time is negative). The negative coefficient on price² suggests nonlinearity in response to price changes and increase in probability of selling as price increases. The positive coefficient on the future expected price—which is the expected price in period 50 given the current price, the number of periods left, and the price generating Markov process—shows that the selling behavior is modulated by future considerations. In particular, a higher expected price in the future makes participants keep the asset longer.

analysis that participants exhibit future regret avoidance, since the significant effect of the variable future expected peak can have other sources. Regressions do not account for the dynamic structure of the optimization problem and essentially just reveal correlations between the selling events and the corresponding states of the market. Therefore, the effect of the future expected peak can come from the attempts of participants to act upon some kind of future expectations, which may not imply that they try to avoid future regret. The same can be said about the possible interaction between past and future regret, the presence of which our data suggest: it can simply be an artifact of optimization with some considerations of the future. In order to resolve this issue, in the following sections, we formulate and provide estimates of the structural model, which allows us to explicitly separate the role of past and future regret from that of future expectations while taking into account the dynamic nature of the task.

III. Regret-Averse Utility Function

We start with defining the regret-averse utility function that is further used in the structural model. We hypothesize that the *highest price in the past*, or *past peak*, defined as $s_{p,t} = \max_{\tau \leq t} y_{\tau}$, is a reference point that our participants use to measure how well they are doing (as shown in Figure 4, panel A). This is a dynamic variable that changes when the price gets above the observed highest peak.⁸ We conjecture that given the current price, which is always less than or equal to $s_{p,t}$, the higher the past peak, the more negative the feeling of *past regret* should be. This implies that if an agent is influenced by past regret, her utility should be negatively proportional to $s_{p,t}$. This dependency, in its turn, influences the decision to keep the asset. This modeling choice is motivated by recent work (e.g., Gneezy 2005; Strack and Viefers 2021) that leverages on the saliency of the highest past price as the key measure of regret, allowing us to disregard other functions of past prices that could be used as reference points for past regret. Highest past prices were found to be important in trading decisions in financial markets. For example, in their analysis of the decision to exercise stock options, Heath, Huddart, and Lang (1999) found that exercising activity doubles when the current price attains the maximum level over the past year.

If participants are aware that they will observe prices even after selling the asset, they can anticipate a situation where the future price will exceed the selling price, which would lead to negative emotions that we call *future regret*. In this case, participants' decisions to sell should be sensitive to the *future expected highest price*, which is a dynamic variable that depends on the current price and the number of periods left before the market closure. When this information is not available, the future regret should not play any role in the selling decisions since participants do not anticipate any negative emotions from observing high prices after selling.⁹ The

⁸Notice that without regret, the optimal policy is to sell the asset whenever the price rises above the threshold in Figure 2 that depends only on the number of periods left. Thus, in the no-regret case, the selling decision is independent of any reference points.

⁹A similar negative response was found in Cooke, Meyvis, and Schwartz (2001), where reported satisfaction scores were negatively correlated with the prices after the sale decision.

expectation of the highest future peak at time t , denoted by $s_{f,t} = E[\max_{t < \tau \leq T} y_\tau | y_t]$, is a function of the price today and the number of periods left until the market closure. The variable $s_{f,t}$ is the expectation of the maximum price achievable in $T - t$ periods given the current price y_t over all possible price paths. An agent sensitive to future regret should have a utility function that is negatively proportional to $s_{f,t}$, but only when the agent knows that the future prices will be observed.

As an additional observation, notice that given a fixed price y_t , $s_{f,t}$ is decreasing in t because of the presence of a terminal period (see online Appendix F for the details). Conversely, $s_{p,t}$ is a weakly increasing function of time, since it is defined as a maximum of the past prices. This suggests that future regret should be dominant in early periods, while past regret should be in late periods. In order to make the utility specification more flexible and to be able to infer whether the current reference point is the highest price observed in the past, the expected highest price in the future, or a combination of these two variables, we add an interaction term to the utility function and specify it as follows:

$$(1) \quad u(y_t, s_{p,t}, s_{f,t}) = \pi y_t - \omega s_{p,t} - \alpha s_{f,t} - \lambda s_{p,t} s_{f,t}.$$

The interaction term can incorporate many types of dependencies between past and future regret. For example, if $\lambda > 0$, then past and future regret are *complements*, one reinforces the other. If $\lambda < 0$, then the two types of regret are *substitutes*, which means that the presence of one type makes the influence of the other one weaker.¹⁰

The consumption part of the utility is given by πy_t , with $\pi \geq 0$.¹¹ The disutility from past regret is captured by the second term with parameter ω , and the disutility from future regret by the third term with parameter α . The parameter λ determines how past and future regret interact. The utility function in (1) offers a simple way to test our predictions: when $\omega = 0$, the decision of the agent does not depend on the past peak or on the future expected peak when $\alpha = \lambda = 0$. This means that the less the agent cares about past/future regret, the less his selling price is influenced by the past/future peak.¹²

Our hypotheses stated above imply that ω should be positive. Note that α and λ should be zero in the No Info condition, since future prices are not available (though participants can, in principle, calculate $s_{f,t}$ in this case and be influenced by it). In the Info condition, α should be positive, while the value of λ in the Info condition can

¹⁰ Alternative specifications of the interaction term are possible. However, we use a parsimonious specification $s_{p,t} s_{f,t}$ because it allows identification of λ and can tell whether past and future regret are complements or substitutes.

¹¹ The results in the previous section show a limited role for risk aversion, so for the analysis reported below, we assume risk neutrality, though we also estimate the model assuming CRRA preferences. When allowing for risk preferences (online Appendix I.2), we define the future regret as the disutility at the future highest peak, $-\alpha U(s_{f,t}; \rho)$, where $U(\cdot; \rho)$ is a CRRA utility with risk aversion coefficient ρ . An alternative approach would be to define it as the expectation of a regret function over possible draws of the future price, e.g., $E[U(\max_{t < \tau \leq T} (y_\tau)) | y_t]$. This, however, would entail significant estimation difficulties. Our definition is in line with the idea that participants have a "target income" at any point in time (see, e.g., Camerer et al. 1997; Crawford and Meng 2011).

¹² It should also be acknowledged that this is not a standard regret aversion function that has one reference point and two parameters like in Bell (1982) and Loomes and Sugden (1982). Since we focus on two reference points (past and future regret), such a function would complicate both the estimation and the interpretation of the results across conditions. Also, we opted for the linear utility, as the nonlinearity of classic utilities with regret would be infeasible to estimate in our setting due to the already complex calculations involved in computing $s_{f,t}$.

be anything depending on the nature of interaction between past and future regret. Thus, estimating the three regret parameters in the two conditions will allow us to test our ideas about the role of past and future regret and, in addition, will make it possible to tell how the reference point changes in time depending on the relative sizes of past and future regret.

Our definition of the future regret is a major departure from the analysis in Strack and Viefers (2021), who focus only on the regret over past decisions. The novelty of our approach is that we consider a decision-maker who takes into account both the endogenously changing past reference point (the past peak) and the exogenously given information about the possibility of future regret, which shares features with the classical static regret. Thus, our decision-maker is affected by both the past price shocks, as in Strack and Viefers (2021), and by the knowledge of the availability of price information after selling the asset, which comes at a cost since the decision-maker may be future regret averse. We model these two forces with separate reference points, one in the past and one in the future, and investigate empirically whether they subdue or reinforce each other.

IV. A Structural Model of Dynamic Regret Avoidance

To assess the role of dynamic regret avoidance in decision-making, we assume that participants follow an optimal policy when choosing to sell the asset given some parameters of the regret-averse utility function. We estimate a dynamic discrete choice model (e.g., Rust 1987, 1994) where the value from selling the asset is directly compared with the continuation value: participants sell when the former is larger than the latter. This section sketches the model that will be taken to the data in Section V. Online Appendix H provides the full derivation.

In our experiment, the evolution of the price of the asset is Markovian, as the price in period $t + 1$ depends only on the price in period t . So participants decide to sell the asset if the current outcome is greater than the discounted value of the future outcomes, which can be represented by a value function. In each period t , one of two choices is made: to sell the asset ($d_t = 0$) or to keep it ($d_t = 1$). As in Section III, $u(x_t)$ denotes the regret-averse utility from selling the asset when the current state is $x_t = (y_t, s_{p,t}, s_{f,t}) \in \mathcal{X}$. A decision-maker's intertemporal expected utility is

$$E \left[\sum_{t=1}^T \beta^{t-1} (1 - d_t) u(x_t) + \varepsilon_t^{d_t} \right],$$

where the expectation is taken over the values of the independent variables x_t and $\beta \in (0, 1)$ is the discounting factor. Similarly to most of the binary static discrete choice models, the value of each choice includes additive iid extreme value type I errors ($\varepsilon_t = (\varepsilon_t^0, \varepsilon_t^1)$), which account for unobserved variables that may affect the decisions. Notice that the decision-maker receives actual utility in only one period when the asset is sold. However, the common structure of the discrete choice models assumes intertemporal optimization with random shocks to the utility, which necessitates the formulation of the expected utility above.

The dynamic environment can be summarized using a value function v_t^d , which represents the time discounted utility obtained by the decision-maker who follows the optimal policy at t :

$$(2) \quad v_t^d(x_t) = \begin{cases} u(x_t), & \text{if } d_t = 0 \text{ (sell);} \\ \beta E_{\mathcal{X}} \left[E_{\varepsilon} \left[\max \left\{ u(x_{t+1}) + \varepsilon_{t+1}^0, v_{t+1}^1(x_{t+1}) + \varepsilon_{t+1}^1 \right\} \right] \middle| x_t \right], & \text{if } d_t = 1 \text{ (keep).} \end{cases}$$

This equation summarizes the decision problem of the agent: given current state x_t , she will keep the asset if this provides more utility than selling it, i.e., if $v_t^1(x_t) > v_t^0(x_t)$. The payoff from keeping the asset corresponds to the discounted value from behaving optimally in the next period. Thus, v^1 includes the expectation over the state variables in the next period, x_{t+1} , and the errors, ε_{t+1} .

The large size of the state space \mathcal{X} and the large number of periods make a solution by backward induction (the classic method when periods are finite) a daunting task. To estimate (2), we rely on the fact that the distribution of the observed choices uniquely identifies the utility function (Hotz and Miller 1993), which allows us to transform (2) into a set of equations that can be estimated by the least squares method (e.g., Pesendorfer and Schmidt-Dengler 2008).

Intuitively, the agent will choose to keep or sell the asset depending on which action provides the higher value conditional on any given realization of the state variable x_t . Therefore, we expect this relation to be reflected in the probability of choosing each action conditional on the state and period. For this case, Hotz and Miller (1993) show the existence of an invertible mapping between the value functions and the related probability of choosing each action given x_t . This probability is known as the conditional choice probability (CCP) and is denoted by $p_t^1(x_t) = \Pr(d_t = 1 | x_t)$ for the probability of continuing and $p_t^0(x_t) = \Pr(d_t = 0 | x_t)$ for the probability of selling at t . Since the CCP can be estimated directly from the data, we treat $p_t^d(x_t)$ as a known object for all t and x_t . The identification procedure uses the CCP—together with the properties of the logit distribution—to express equation (2) in terms of data.

The logit assumption gives an analytical solution for the probability of choosing each action. For example, the probability of selling is $p_t^0(x_t) = 1 / (1 + \exp(v_t^1(x_t) - v_t^0(x_t)))$, which depends on the difference between the values of keeping and selling in (2). This difference is

$$(3) \quad v_t^1(x_t) - v_t^0(x_t) \\ = -u(x_t) + \beta E_{\mathcal{X}} \left[E_{\varepsilon} \left[\max \left\{ v_t^0(x_{t+1}) + \varepsilon_{t+1}^0, v_t^1(x_{t+1}) + \varepsilon_{t+1}^1 \right\} \right] \middle| x_t \right].$$

This equation can be simplified further by exploiting the properties of the logistic error structure, as in Hotz and Miller (1993). First, let's consider the LHS of (3). We can rewrite $v_t^1(x_t) - v_t^0(x_t)$ as a function of $p_t^0(x_t)$ using their relationship shown

above. Denote this function by $\phi(p_t^0(x_t)) \equiv \ln(1 - p_t^0(x_t)) - \ln(p_t^0(x_t))$.¹³ Thus, the difference between the value of keeping and selling the asset can be thought of in terms of changes in the probability of selling the asset. This means that the left-hand side of (3) is a known function of the data, the CCP. Next, the inner expectation and the max operator in the right-hand side of (3) can also be simplified using the properties of the logit errors. Online Appendix H shows all the steps of this derivation. As a result, (3) becomes

$$(4) \quad \phi(p_t^0(x_t)) = -u(x_t) + \beta \sum_{x_{t+1} \in \mathcal{X}} \left(u(x_{t+1}) - \ln(p_{t+1}^0(x_{t+1})) \right) f(x_{t+1} | x_t),$$

where $f(x_{t+1} | x_t)$ is the known transition probability between consecutive periods in \mathcal{X} estimated from the data, and the summation substitutes the integration with respect to x_{t+1} as we discretize the state space before estimation.¹⁴

Several observations about the equation (4) should be made. First, it summarizes intertemporal choices by only comparing the gain from selling in the next period (i.e., $\sum_{x_{t+1} \in \mathcal{X}} u(x_{t+1}) f(x_{t+1} | x_t)$) with the expected (log) probability of selling in the next period given by $-\sum_{x_{t+1} \in \mathcal{X}} \ln(p_{t+1}^0(x_{t+1})) f(x_{t+1} | x_t)$. This last term is important because it incorporates the continuation value and can be thought of as the utility from waiting for a better price. In fact, this expectation is proportional to the continuation value at $t + 1$ through the definition of the CCP.¹⁵ Hence, we know that the right-hand side of (4) increases when agents expect high returns from keeping the asset in the following periods. Because the left-hand side of (4) corresponds to the difference between the value functions from keeping and selling the asset, a greater continuation value implies that the agent is more likely to keep the asset in period t .

Second, given the same continuation value, if $\omega > 0$, the model predicts that the agent will be less likely to sell in period t if the distance between the past peak and the current price is larger than the expectation of the same difference in the following period. To see this, notice that the right-hand side of (4) increases if the difference between the past peak and the current price goes up, which in turn should increase the left-hand side or decrease the probability of selling. This reasoning can also be applied to the future expected peak. In the Info condition, the agent will be less likely to sell the asset in period t if $s_{f,t} > E[s_{f,t+1}]$ and $\alpha > 0$, other things equal. This follows again from the increase in the right-hand side of (4) when the expectation of the future peak changes marginally.

We have constructed a simple two-step estimator. The first step involves recovering the CCP and the transition matrix directly from the data. In the second step, these objects are plugged into (3). This gives us the objective function (equation (4)) used

¹³The CCP of selling the asset is $p_{t+1}^0(x_t) = 1/(1 + \exp(v_t^1(x_t) - v_t^0(x_t)))$. This can be transformed into $v_t^1(x_t) - v_t^0(x_t) = \ln(1 - p_t^0(x_t)) - \ln(p_t^0(x_t))$.

¹⁴The discretization of the state space is necessary to estimate the model. For our experiment, this is not a problem; the participants face a discrete state space anyway, as y_t was rounded to cents. The discretization is implemented according to the approach proposed by Tauchen (1986) to approximate a vector autoregression model with a finite state Markov chain. All variables (current price, past peak, and future peak) are discretized on the same support in $[0.59, 9.32]$. The distance between the 400 bins is €0.02. This method is described in detail in online Appendix G.

¹⁵From the derivations above, we have $-\ln(p_{t+1}^0(x_{t+1})) = \ln(1 + \exp(v_{t+1}^1(x_{t+1}) - v_{t+1}^0(x_{t+1})))$, which is approximately equal to $v_{t+1}^1(x_{t+1}) - v_{t+1}^0(x_{t+1})$.

to estimate a parameterized version of the utility of selling the asset $u(x_t)$, which includes regret-averse components for the two conditions (Info and No Info). In conclusion, the procedure just described relies on the common logit assumption in the binary choice literature, the presence of a terminating action (selling the asset), and the Markovian nature of the changes in the state variables.

V. Estimation of the Structural Model

We now turn to the estimation of the dynamic discrete choice model in Section IV. However, before proceeding to the estimation of (3), we analyze how the CCP differs in the two conditions, as this can further elucidate the mechanisms at play.

A. Estimation of the Conditional Choice Probabilities

The conditional probability of selling the asset (or continuing) at period t is computed directly from the data. We exclude periods 15 and 50 since no one sold the asset in the former (first choice period) and the choice is forced in the latter (last period). Participants sell their asset in different periods, resulting in a highly unbalanced dataset. The CCPs are constructed using a logit estimator of the choices of the active participants in each period $t \in \{16, \dots, 49\}$ as a function of the realized state variables. It is important to stress that there are two policy functions to be estimated for each period since the experiment has two conditions. The CCP for either the No Info or Info condition in period t can be represented as follows:

$$(5) \quad \Pr\{d_t = 0 | x_t\} = \Lambda(\beta_{1t}y_t + \beta_{2t}s_{p,t} + \beta_{3t}s_{f,t}), \quad \forall t \in \{16, \dots, 49\},$$

where $\Lambda(\cdot)$ stands for the logistic distribution. In principle, several other valid specifications can be used. However, since the sample size shrinks as participants sell their assets over time, adding additional covariates undermines the identification of the parameters.¹⁶

Figure 5 shows the projections of the time-averaged fitted CCP in the No Info condition.¹⁷ Specifically, each line represents the estimate of the probability of selling that results from averaging the fitted values of 34 logit regressions (1 for each time period).¹⁸ For prices below €5, the probability of selling is the highest when the past peak is €3, is lower when the past peak is €5, and is close to 0 for past peaks €7 and €8. This means that when prices are low, the participants are strongly influenced by the size of the past peak and wait for the price to become closer to it. For the past peaks €7 and €8, which are very common in our data, the probability of

¹⁶ Adding square and interaction terms creates a large multicollinearity problem, eventually impairing the identification of the β_{it} coefficients. In fact, the singular value decomposition of the matrices of covariates in (5) show that including these terms makes it ill conditioned in most periods. Also, clustering at subject level does not affect the results.

¹⁷ For the purpose of making this graph illustrative, the CCPs in Figure 5 were calculated without the future regret term, or assuming $\beta_{3t} = 0$ in (5).

¹⁸ Thus, the CCPs in Figures 5 and 6 are shown just for illustration. They are out-of-sample estimates that do not take into account the influence of the current price on the past peak (i.e., the current price cannot be larger than the highest observed price).




FIGURE 5. PROBABILITY OF SELLING THE ASSET IN THE NO INFO CONDITION

Notes: The effect of the past peak on the probability of selling the asset in the No Info condition. The CCP is computed by taking the average of the fitted values for all periods $t \in \{16, \dots, 49\}$. This figure is for illustrative purpose only and is based on a state space discretized over 200 bins.

selling increases rapidly when the price approaches €7. This demonstrates that the past peak indeed serves as a reference point.

Figure 6 illustrates similar projections of the CCP in the Info condition. For fixed value of future regret, the relationship between the curves with past regret equal to €5 and €7 is the same as in Figure 5. However, the effect of past regret is much smaller in this case. We conjecture that this is due to the presence of the future regret term that dominates the past regret. In what follows, we show that there is a substitution effect between past and future regret that can explain this pattern.

B. Estimation of the Parameters

In order to causally connect regret avoidance and decisions to sell in our experiment, we estimate (4) by nonlinear least squares procedure (Bajari et al. 2016).

We proceed with the estimation of a parametric version of (4) where the per-period utility from selling is defined as

$$(6) \quad u(y_t, s_{p,t}, s_{f,t}) = \pi y_t - \mathbf{1}_{No\ info}(\omega_{NI}s_{p,t} + \alpha_{NI}s_{f,t} + \lambda_{NI}s_{p,t}s_{f,t}) \\ - \mathbf{1}_{Info}(\omega_I s_{p,t} + \alpha_I s_{f,t} + \lambda_I s_{p,t}s_{f,t}).$$




FIGURE 6. PROBABILITY OF SELLING THE ASSET IN THE INFO CONDITION

Notes: The effect of the past peak and the expected future peak in the Info condition. The CCP is computed by taking the average of the fitted values from (5) for all periods $t \in \{16, \dots, 49\}$. This figure is for illustrative purpose only and is based on a state space discretized over 200 bins.

This is the regret-averse utility discussed in Section III. Specifically, we assume that participants take into account both past and future regret in both the Info and No Info conditions but do it in potentially different ways. The estimation of the six regret parameters denoted by the subscripts “NI” for No Info and “I” for Info conditions should provide support to our hypotheses.¹⁹

Table 2 shows the estimated parameters of the utility function. As we have hypothesized, the estimates of the past regret parameters $\hat{\omega}_{NI}$ and $\hat{\omega}_I$ are positive and significant for all discount factor specifications. This demonstrates that our participants respond to past regret in both conditions. We also observe that $\hat{\alpha}_I$ is positive and significant, which implies that participants have future regret in the markets where they know that future prices will be available. Importantly, the coefficients $\hat{\alpha}_{NI}$ and $\hat{\lambda}_{NI}$ are *not* significant in all models. Thus, we conclude that the future expected peak plays no role in the decisions to sell when the participants know that they will not observe the future prices after selling. The observation that future expected peak only enters selling decisions when future prices are observable is crucial since it refutes any theory that incorporates $s_{f,t}$ in the utility function but is not based on emotional reactions related to observability of future prices. Any such theory would predict no difference between the Info and No Info conditions since participants can easily use $s_{f,t}$ —which is essentially an expectation—for their

¹⁹The indicator functions $\mathbf{1}$ distinguish the utility derived in one condition from the other. The parameters $\pi, \omega_{NI}, \alpha_{NI}, \lambda_{NI}, \omega_I, \alpha_I, \lambda_I$ are free to vary and indicate how strongly participants’ decisions are affected by regret.

TABLE 2—THE ESTIMATION OF (4) IN PERIODS 16 TO 48 FOR DIFFERENT VALUES OF β

Parameter	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
$\hat{\pi}$	1.789 (0.014)	1.788 (0.014)	1.787 (0.014)
$\hat{\omega}_{NI}$	1.432 (0.462)	1.555 (0.413)	1.643 (0.376)
$\hat{\omega}_I$	2.609 (0.473)	2.585 (0.424)	2.562 (0.385)
$\hat{\alpha}_{NI}$	0.134 (0.341)	0.229 (0.303)	0.296 (0.274)
$\hat{\alpha}_I$	1.761 (0.348)	1.719 (0.309)	1.679 (0.281)
$\hat{\lambda}_{NI}$	-0.046 (0.051)	-0.059 (0.046)	-0.068 (0.043)
$\hat{\lambda}_I$	-0.265 (0.053)	-0.260 (0.048)	-0.256 (0.044)
Observations	111,613	111,613	111,613

Notes: Standard errors are in parentheses. The CCP is computed using the formula in (5) for both conditions.

decisions in both conditions. Thus, our findings provide a strong and direct evidence of past and future regret avoidance.

Next, we turn to the interpretation of the coefficient $\hat{\lambda}_I$ on the interaction of past and future regret in the Info condition. Notice that it is negative and significant. This confirms the presence of a *substitution effect* between the two types of regret. The size of $\hat{\lambda}_I$ allows us to conclude that participants are only affected by one type of regret at a time. In particular, they pay attention only to the largest of the two: when either past or future regret is large and the other is small, the interaction term offsets the effect of the small term (see Figure 7 in Section VI). Moreover, the presence of the interaction term implies that participants switch their focus between past and future regret *dynamically* within each market depending on which peak is larger. This suggests that people can be surprisingly flexible at being past or future oriented when it comes to selling decisions in dynamic settings.²⁰

Finally, we verify that our results cannot be explained by loss aversion. A loss occurs if the asset is sold at a price below the purchase price in period 15. Before, in Section II, we have provided arguments that loss aversion cannot explain our data. Here, we go further and explicitly estimate a structural model with utility that has a loss aversion term in it. The estimation is reported in online Appendix I.3. The loss aversion term is not significant. This also supports our results in Table 2: the estimates of the past and future regret stay unchanged. We conclude that loss aversion plays no role in the decision to sell the asset.

²⁰In online Appendix I.1, we also estimate a model in which we assume that the decision-maker does not anticipate future regret in the No Info condition as well as several other robustness analyses (online Appendix I.2). The estimates confirm the conclusions from Table 2.

VI. Discussion

We find a strong imprint of past regret on the decisions of our participants in an optimal stopping experiment. Our main findings, however, lie in the domain of future regret and its dynamic interaction with past regret, and they can be summarized as follows. First, the participants *are able* to contemplate the counterfactual situation in which they sell the asset today and later regret it when the price goes up. Moreover, they take this possibility into account by trying to sell the asset at a price closer to the future expected maximum. Second, the participants are not *always* influenced by future regret. They take it into account only when they know that the information about future prices will be available after they sell the asset. Third, past and future regret do not work independently. They *interact* by offsetting each other, which leads to only the strongest being reflected in the decisions.

When comparing the selling behavior in the No Info and Info conditions, it is important to note that the conditions differ only in the information provided *after* the choice was made. Before the choice, the exactly identical information is conveyed to the decision-maker. Therefore, in principle, it is possible to choose in the same way in both conditions. Namely, nothing stops the participants from calculating the expected future maximum value and acting upon it even if the future prices are not revealed. However, as the estimation of the structural model demonstrates, this is not the case, and the *same participant* who avoids future regret in the Info condition chooses to ignore it in the No Info condition. This is particularly surprising given that making optimal selling decisions in our dynamic environment involves calculating future expected prices *even without deliberation on future regret*. This exposes the complexity of intertemporal choice by the regret-averse participants and, particularly, its sensitivity to the context and information available in the future.²¹

The estimation of the structural model shows a significant interaction effect between past and future regret in the Info condition. Specifically, this interaction is negative and thus works to counteract the effect of the smaller regret term (past or future). This mechanism, though static in nature, creates a compelling *dynamic effect*: the impact of the past and the future on the probability of selling changes in time as the past and future regret terms change in relative size. Figure 7 provides a graphical intuition. In the left graph before period 18, the past regret term, which is dominated by the future regret term, is offset by the interaction. After this period, the roles of the past and future regret terms switch, and the future regret is now offset by the interaction term. Overall, the interaction term in both graphs is close to the *minimum* of the past and future regret terms that makes the higher regret term exert most of the influence on the decision to sell. The participants try to minimize the distance from a *global* highest peak or $\max\{s_{p,t}, s_{f,t}\}$, thus treating the past and the (expected) future in the same way. It should be emphasized that this result has emerged endogenously without introducing the maximum of the two peaks as the definition of the regret function. This also explains the different

²¹The ability to contemplate hypothetical counterfactual scenarios is also experimentally investigated by Esponda and Vespa (2014) in a different environment with multiple agents with strategic interactions and sequential decisions.




FIGURE 7. DYNAMICS OF PAST AND FUTURE REGRET

Notes: Examples of the dynamics of past and future regret in two selected markets in the Info condition for the periods $t \in \{16, \dots, 48\}$. The curves show the terms of the estimated regret function (i.e., past regret = $\hat{\omega}_I s_{p,t}$, future regret = $\hat{\alpha}_I s_{f,t}$, and interaction term = $|\hat{\lambda}_I| s_{f,t} s_{p,t}$) in column 4 of Tabel I4 in online Appendix I. The solid vertical line shows the moment at which the participants switch the focus from future regret to past regret.

rates at which participants in the No Info and Info conditions sell when the current price is in the vicinity of the past peak, as documented in panel B of Figure 4. In the early periods, future regret is a reference point for the participants in the Info condition but not the participants in the No Info condition. As time goes by, the saliency of past regret increases, eventually dominating the future regret term (see Figure 7).

In our experiment, this effect is detected *within subjects*, which means that orientation toward the past or the future can change rapidly depending on the circumstances. More importantly, this implies that the behavior on financial markets can potentially be influenced by seemingly unrelated events that nevertheless refocus the attention of the investors on the past or expected future developments (e.g., Klibanoff, Lamont, and Wizman 1998; Bordalo, Gennaioli, and Shleifer 2018). For example, in our setting, the value of the expected future maximum depends on the number of periods left before the market closure: for any fixed current price, the closer the end, the lower the expected future maximum. Therefore, sudden news that the closure will happen earlier should decrease future regret and thus make investors more wary of the past. This can potentially lead to two outcomes: if the past peak was high and was *dominating* the expected future peak, then nothing should change; however, if the past peak was low and was *dominated* by the expected future peak, then early closure can lead to a selling spree since the dominating regret term—in this case, future regret—has decreased. A similar pattern to the dynamic substitution we elicited in our study was also found across New York taxi drivers in their labor supply decisions (Crawford and Meng 2011). While drivers have flexible schedules and can stop driving after any trip, their choices seem to target either income or

hours worked. In particular, it is the furthest (from the current state) among the two objectives that is the dominant reference point.

The findings of our study add to the existing literature on the multiplicity of reference points (e.g., Kahneman 1992; Baucells, Weber, and Welfens 2011) and their endogenous formation (e.g., Kőszegi and Rabin 2006, 2007; Gill and Prowse 2012) by fully spelling out their mechanism and estimating their relationship in a dynamic setting. We conclude that ex post information shapes agents' actions in our dynamic setting and that agents make no attempt to integrate competing/different reference points but rather dynamically select the most relevant one.

Our results imply another interesting behavioral effect that is concerned with the potential choice between observing and not observing the future price after selling the asset. In particular, the estimates of the utility parameters suggest that having no information should be preferable to having it ($\hat{\omega}_{NI} < \hat{\alpha}_I < \hat{\omega}_I$). So it is not inconceivable that the investors would be willing to pay for not being able to observe the future prices of the asset (e.g., Bell 1983; Caplin and Leahy 2001). This can have consequences for policies directed at the regulation of stock market trading such as short selling (selling to subsequently repurchase an asset), which could be welfare improving over bans (Beber and Pagano 2013). Nevertheless, we would like to stress that the relative size of past and future regret and their interaction is an empirical question that requires case-by-case analysis. Moreover, we believe that our approach could be used to investigate the role of regret avoidance in real-life dynamic situations.

VII. Conclusion

In an experimental task that resembles a stock market, we study how past and future regret avoidance influences selling decisions. We use a dynamic discrete choice model to evaluate the parameters of a utility function that incorporates regret avoidance preferences and find that both past and future regret play an important role in the choices to sell. When participants in the experiment know that after they sell the asset, they will no longer see the evolution of the price, their decisions to sell are strongly influenced by past regret avoidance. Namely, participants keep the asset longer in order to sell at a price close to the highest past price observed. When participants are aware that after they sell the asset, they will continue to observe the price on the market, their choices to sell change: now future regret avoidance also becomes important. Participants take into account the anticipated future regret that they would experience if the price of the asset increased after they sold it, and they try to minimize this effect.

Moreover, we find that past and future regret avoidance do not just influence the decisions in a simple additive way. They interact with each other. In particular, participants pay more attention to the type of regret that is more prominent: if the past highest peak looms higher than the expected future peak, then past regret avoidance enters the decision to sell. If the anticipated regret in the future is larger than the potential past regret, then future regret avoidance becomes important. This substitution effect was not previously mentioned in the literature and may be of particular interest to policymakers.

REFERENCES

- Bajari, Patrick, Chenghuan Sean Chu, Denis Nekipelov, and Minjung Park.** 2016. "Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action." <https://escholarship.org/content/qt9c35w8sq/qt9c35w8sq.pdf?t=ojx49q>.
- Baucells, Manel, Martin Weber, and Frank Welfens.** 2011. "Reference-Point Formation and Updating." *Management Science* 57 (3): 506–19.
- Beber, Alessandro, and Marco Pagano.** 2013. "Short-Selling Bans around the World: Evidence from the 2007–09 Crisis." *Journal of Finance* 68 (1): 343–81.
- Bell, David E.** 1982. "Regret in Decision-Making under Uncertainty." *Operations Research* 30 (5): 961–81.
- Bell, David E.** 1983. "Risk Premiums for Decision Regret." *Management Science* 29 (10): 1156–66.
- Bikhchandani, Sushil, and Uzi Segal.** 2014. "Transitive Regret over Statistically Independent Lotteries." *Journal of Economic Theory* 152: 237–48.
- Bleichrodt, Han, Alessandra Cillo, and Enrico Diecidue.** 2010. "A Quantitative Measurement of Regret Theory." *Management Science* 56 (1): 161–75.
- Boeri, Marco, Alberto Longo, José M. Grisolia, W. George Hutchinson, and Frank Kee.** 2013. "The Role of Regret Minimisation in Lifestyle Choices Affecting the Risk of Coronary Heart Disease." *Journal of Health Economics* 32 (1): 253–260.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer.** 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1): 199–277.
- Buturak, Gökhan, and Özgür Evren.** 2017. "Choice Overload and Asymmetric Regret." *Theoretical Economics* 12 (3): 1029–56.
- Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler.** 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics* 112 (2): 407–41.
- Camille, Nathalie, Giorgio Coricelli, Jerome Sallet, Pascale Pradat-Diehl, Jean-René Duhamel, and Angela Sirigu.** 2004. "The Involvement of the Orbitofrontal Cortex in the Experience of Regret." *Science* 304 (5674): 1167–70.
- Caplin, Andrew, and John Leahy.** 2001. "Psychological Expected Utility Theory and Anticipatory Feelings." *Quarterly Journal of Economics* 116 (1): 55–79.
- Cooke, Alan D.J., Tom Meyvis, and Alan Schwartz.** 2001. "Avoiding Future Regret in Purchase-Timing Decisions." *Journal of Consumer Research* 27 (4): 447–59.
- Coricelli, Giorgio, Hugo D. Critchley, Mateus Joffily, John P. O'Doherty, Angela Sirigu, and Raymond J. Dolan.** 2005. "Regret and Its Avoidance: A Neuroimaging Study of Choice Behavior." *Nature Neuroscience* 8 (9): 1255–62.
- Crawford, Vincent P., and Juanjuan Meng.** 2011. "New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational Expectations Targets for Hours and Income." *American Economic Review* 101 (5): 1912–32.
- Esponda, Ignacio, and Emanuel Vespa.** 2014. "Hypothetical Thinking and Information Extraction in the Laboratory." *American Economic Journal: Microeconomics* 6 (4): 180–202.
- Filiz-Ozbay, Emel, and Erkut Y. Ozbay.** 2007. "Auctions with Anticipated Regret: Theory and Experiment." *American Economic Review* 97 (4): 1407–18.
- Fioretti, Michele, Alexander Vostroknutov, and Giorgio Coricelli.** 2022. "Replication data for: Dynamic Regret Avoidance." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.38886/E130441V1>.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics* 10 (2): 171–78.
- Fogel, Suzanne O'Curry, and Thomas Berry.** 2006. "The Disposition Effect and Individual Investor Decisions: The Roles of Regret and Counterfactual Alternatives." *Journal of Behavioral Finance* 7 (2): 107–16.
- Frydman, Cary, and Colin Camerer.** 2016. "Neural Evidence of Regret and Its Implications for Investor Behavior." *Review of Financial Studies* 29 (11): 3108–39.
- Frydman, Cary, Samuel M. Hartzmark, and David H. Solomon.** 2018. "Rolling Mental Accounts." *Review of Financial Studies* 31 (1): 362–97.
- Gill, David, and Victoria Prowse.** 2012. "A Structural Analysis of Disappointment Aversion in a Real Effort Competition." *American Economic Review* 102 (1): 469–503.
- Gneezy, Uri.** 2005. "Updating the Reference Level: Experimental Evidence." In *Experimental Business Research*, Vol. 3, edited by Rami Zwick and Amnon Rapoport, 263–84. Boston, MA: Springer.

- Hayashi, Takashi.** 2008. "Regret Aversion and Opportunity Dependence." *Journal of Economic Theory* 139 (1): 242–68.
- Hayashi, Takashi, and Hisayuki Yoshimoto.** 2016. "Risk- and Regret-Averse Bidders in Sealed-Bid Auctions." Unpublished.
- Hazan, Elad, and Satyen Kale.** 2015. "An Online Portfolio Selection Algorithm with Regret Logarithmic in Price Variation." *Mathematical Finance* 25 (2): 288–310.
- Heath, Chip, Steven Huddart, and Mark Lang.** 1999. "Psychological Factors and Stock Option Exercise." *Quarterly Journal of Economics* 114 (2): 601–27.
- Holt, Charles A., and Susan K. Laury.** 2002. "Risk Aversion and Incentive Effects." *American Economic Review* 92 (5): 1644–55.
- Hotz, V. Joseph, and Robert A. Miller.** 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *Review of Economic Studies* 60 (3): 497–529.
- Kahneman, Daniel.** 1992. "Reference Points, Anchors, Norms, and Mixed Feelings." *Organizational Behavior and Human Decision Processes* 51 (2): 296–312.
- Klibanoff, Peter, Owen Lamont, and Thierry A. Wizman.** 1998. "Investor Reaction to Salient News in Closed-End Country Funds." *Journal of Finance* 53 (2): 673–99.
- Kőszegi, Botond, and Matthew Rabin.** 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.
- Kőszegi, Botond, and Matthew Rabin.** 2007. "Reference-Dependent Risk Attitudes." *American Economic Review* 97 (4): 1047–73.
- Leung, Samantha, and Joseph Y. Halpern.** 2015. "Minimizing Regret in Dynamic Decision Problems." <https://arxiv.org/pdf/1502.00152.pdf>.
- Loomes, Graham, and Robert Sugden.** 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Economic Journal* 92 (368): 805–24.
- Michenaud, Sébastien, and Bruno Solnik.** 2008. "Applying Regret Theory to Investment Choices: Currency Hedging Decisions." *Journal of International Money and Finance* 27 (5): 677–94.
- Muermann, Alexander, Olivia S. Mitchell, and Jacqueline M. Volkman.** 2006. "Regret, Portfolio Choice, and Guarantees in Defined Contribution Schemes." *Insurance: Mathematics and Economics* 39 (2): 219–29.
- Oprea, Ryan.** 2014. "Survival versus Profit Maximization in a Dynamic Stochastic Experiment." *Econometrica* 82 (6): 2225–55.
- Oprea, Ryan, Daniel Friedman, and Steven T. Anderson.** 2009. "Learning to Wait: a Laboratory Investigation." *Review of Economic Studies* 76 (3): 1103–24.
- Pesendorfer, Martin, and Philipp Schmidt-Dengler.** 2008. "Asymptotic least squares estimators for dynamic games." *Review of Economic Studies* 75 (3): 901–28.
- Qin, Jie.** 2015. "A Model of Regret, Investor Behavior, and Market Turbulence." *Journal of Economic Theory* 160: 150–74.
- Rust, John.** 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033.
- Rust, John.** 1994. "Structural Estimation of Markov Decision Processes." In *Handbook of Econometrics*, Vol. 4, edited by Robert F. Engle and Daniel L. McFadden, 3081–3143. Amsterdam: North Holland.
- Sarver, Todd.** 2008. "Anticipating Regret: Why Fewer Options May Be Better." *Econometrica* 76 (2): 263–305.
- Skiadas, Costis.** 1997. "Conditioning and Aggregation of Preferences." *Econometrica* 65 (2): 347–67.
- Strack, Philipp, and Paul Viefers.** 2021. "Too Proud to Stop: Regret in Dynamic Decisions." *Journal of the European Economic Association*, 19 (1): 165–99.
- Tauchen, George.** 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20 (2): 177–81.
- Zeelenberg, Marcel.** 1999. "Anticipated Regret, Expected Feedback and Behavioral Decision-Making." *Journal of Behavioral Decision-Making* 12 (2): 93.

Dynamic Regret Avoidance

Fioretti, Vostroknutov, Coricelli

Online Appendix

A Experimental Design

In the experiment participants made choices in 48 “stock markets,” presented to each of them in individually generated random order. In each market a participant was shown the price dynamics unfolding in real time either until the asset was sold or until market closure after 50 periods. The price updated each 0.8 seconds. First, participants observed the market price evolve for 15 periods. Then they “entered” the market. In the instructions this was presented as if they bought an asset in period 15. After this, participants kept observing the evolution of the market price and had to decide when to “sell the asset.” The payoff, or profit, that each participant received in each market was equal to the selling price minus the entry price. Participants were paid for only one randomly chosen market. No one could lose money if the profit of the chosen market was negative, since participants were given an initial endowment of €10 that covered the highest possible loss.

Each participant was making choices in two types of markets, which differed only in the amount of information that participants received *after* they have sold the asset. In the No Info condition, after selling the asset, no information about the future evolution of the price was provided. In the Info condition, after selling the asset, participants observed how the price changed until the end of that market. In both cases the participants could not change their decision after they have sold the asset. The market condition (No Info or Info) was shown from period 1 on in the upper-left corner of the market graph (see figures below).

Overall, 154 participants took part in the experiment. All sessions were run in March 2017 at the CEEL laboratory, Department of Economics, University of Trento. Another set of 135 participants took part in the experiment in June 2016 in the same laboratory. These data are not reported in this paper. In the June 2016 experiment participants were not informed about the process that generated the price and were not given initial training (see below). Otherwise the two experiments were identical. One session in the June 2016 experiment was aborted due to the network overload and the data was discarded. The data for one participant in the June 2016 experiment was discarded, as she had to leave the experiment in the middle of the market task. No other sessions or pilots were conducted. The experiments were programmed in z-Tree (Fischbacher, 2007).

A.1 Market Details

The price dynamics for each market was generated randomly using the process $y_{t+1} = \alpha y_t + (1 - \alpha)\varepsilon$, where y_{t+1} is the price in period $t + 1$, $\alpha = 0.6$ is a fixed constant for all markets and $\varepsilon \sim U[0, 10]$ is an iid random variable (uniform on $[0, 10]$). In period 1 each market started from price €2.5, €5, or €7.5. Thus, the price changed in the range from €0 to €10. All participants saw the same price dynamics for a given market. Each market lasted for 50 periods, which was known to the participants. In period 15 of each market the participants were forced to enter the market. This was explained to them in the instructions in terms of their buying an object on the market in period 15 for the current market price (see instructions in Appendix J). Then the participants were instructed that they can sell the asset at any time before period 50 and that their earnings in that market would be equal to the difference between the selling price and the entry price (if they did not sell their earnings were equal to the price in period 50 minus the price in period 15). The prices on the market were presented in actual Euros, so no tokens were used and there was no need for having an exchange rate. All the information about the current market condition, the entry price, the selling price and the current price was presented on the screen at appropriate times. Descriptions under Figures A1 and A2 explain.




Figure A1: The left picture shows the market price evolution before period 15, which is marked by a vertical red line. At period 15 the market “stopped,” so that participants could inspect the entry price. An ENTER (ENTRATA) button should have been pressed to start the market again. After period 15 the participants could check the entry price by looking at the top left of the screen where it was indicated in red (right picture). To sell the asset participants needed to press EXIT (USCITA) button.



Figure A2: The right picture shows the market in Info condition *after* a participant sold the asset (the period of selling is indicated by a blue vertical line). After selling the asset, the participant could see the selling price in blue and the profit in green or red, depending on whether the profit was positive or negative (on top of the screen). In addition, the participant observed the future evolution of the price until period 50 (the price changed each 0.8 seconds). In the No Info condition (left picture) everything was the same except that the participant did not observe the future price, but still had to wait until the market closure. The sentence at the bottom of the left picture says: “Please wait until the market is closed.”

The timing of each market was as following. The new price was shown each 0.8 seconds.¹ This was long enough for participants to be able to react and sell the asset at the current price if they chose to do so. In the Info condition participants had to observe the evolution of the price until period 50: they could not skip to the next market. In the No Info treatment, after selling the asset, they had to wait until the market reached period 50 (without observing the price). This was done in order to 1) remove the incentive to go quicker through the task and 2) make No Info and Info conditions as similar as possible.

¹The experiment was implemented in z-Tree (Fischbacher, 2007), which does not allow for precise time control. Thus, the actual time between periods could have been slightly larger.

A.2 Price Dynamics and Training

Participants were explicitly informed about the process that generates the price (see instructions in Appendix J). The formula $y_{t+1} = \alpha y_t + (1 - \alpha)\varepsilon$ was explained to them and four examples of the price range in the next period depending on the current price were given.

Participants went through a series of six training markets which could not be chosen for the payment. The training markets were in all respects identical to the real markets except the phrase ROUND DI PROVA (“training round”) written across the background in a very large font. Out of six training markets two started at €2.5, two at €5, and two at €7.5. One market in each pair was presented in the No Info and one in the Info condition. The sequence of markets and conditions were independently randomized individually for each participant.

A.3 Overall Design Details

Participants chose in 48 markets. The price dynamics for each market was pre-generated using the rule described above (see Figure A3 below). Thus, each participant chose in exactly the same markets. For the three subsets of 16 markets the starting price was equal to €2.5, €5, or €7.5. The order of the markets was randomized in real time for each participant. Thus, there is only an infinitesimal probability that any two participants saw the same sequence of markets. The market condition, No Info or Info, was determined in the following way. Half of the 16 markets of each kind (starting at €2.5, €5, €7.5) were randomly assigned to the condition No Info and another half to the condition Info. Thus, equal number of markets of each of the three kinds were shown in the two conditions. The participants assigned to the computers with odd numbers saw markets in these predetermined conditions. The participants assigned to the computers with even numbers saw the same markets in the opposite conditions. Thus, for each given market, there is an (approximately) equal number of participants who saw that market in the No Info and Info conditions.

When participants sold the asset they could see their profit (see Figure A2). However, the participants were informed that they will be paid for only one randomly chosen market. In order to avoid losses, the participants were given €10 at the beginning of the experiment, so their earnings after the market task were €10 plus the profit in one randomly chosen market (which could have been negative).

A.4 Additional Tasks

After choosing in the sequence of 48 markets the participants were presented with the Holt and Laury task (Holt and Laury, 2002). We did not use the original payoffs from Holt and Laury (2002) as our participants could have seen those before. Instead we took the equivalent payoffs from Eijkelenboom *et al.* (2016). The instructions and the screenshots are presented in Appendix K.2. The participants, in addition to their earnings in the market task, received the payoff from one of the lotteries that they chose in the Holt and Laury task.

In the end of the experiment the participants were given a sequence of standard demographic questions.

A.5 Market Prices

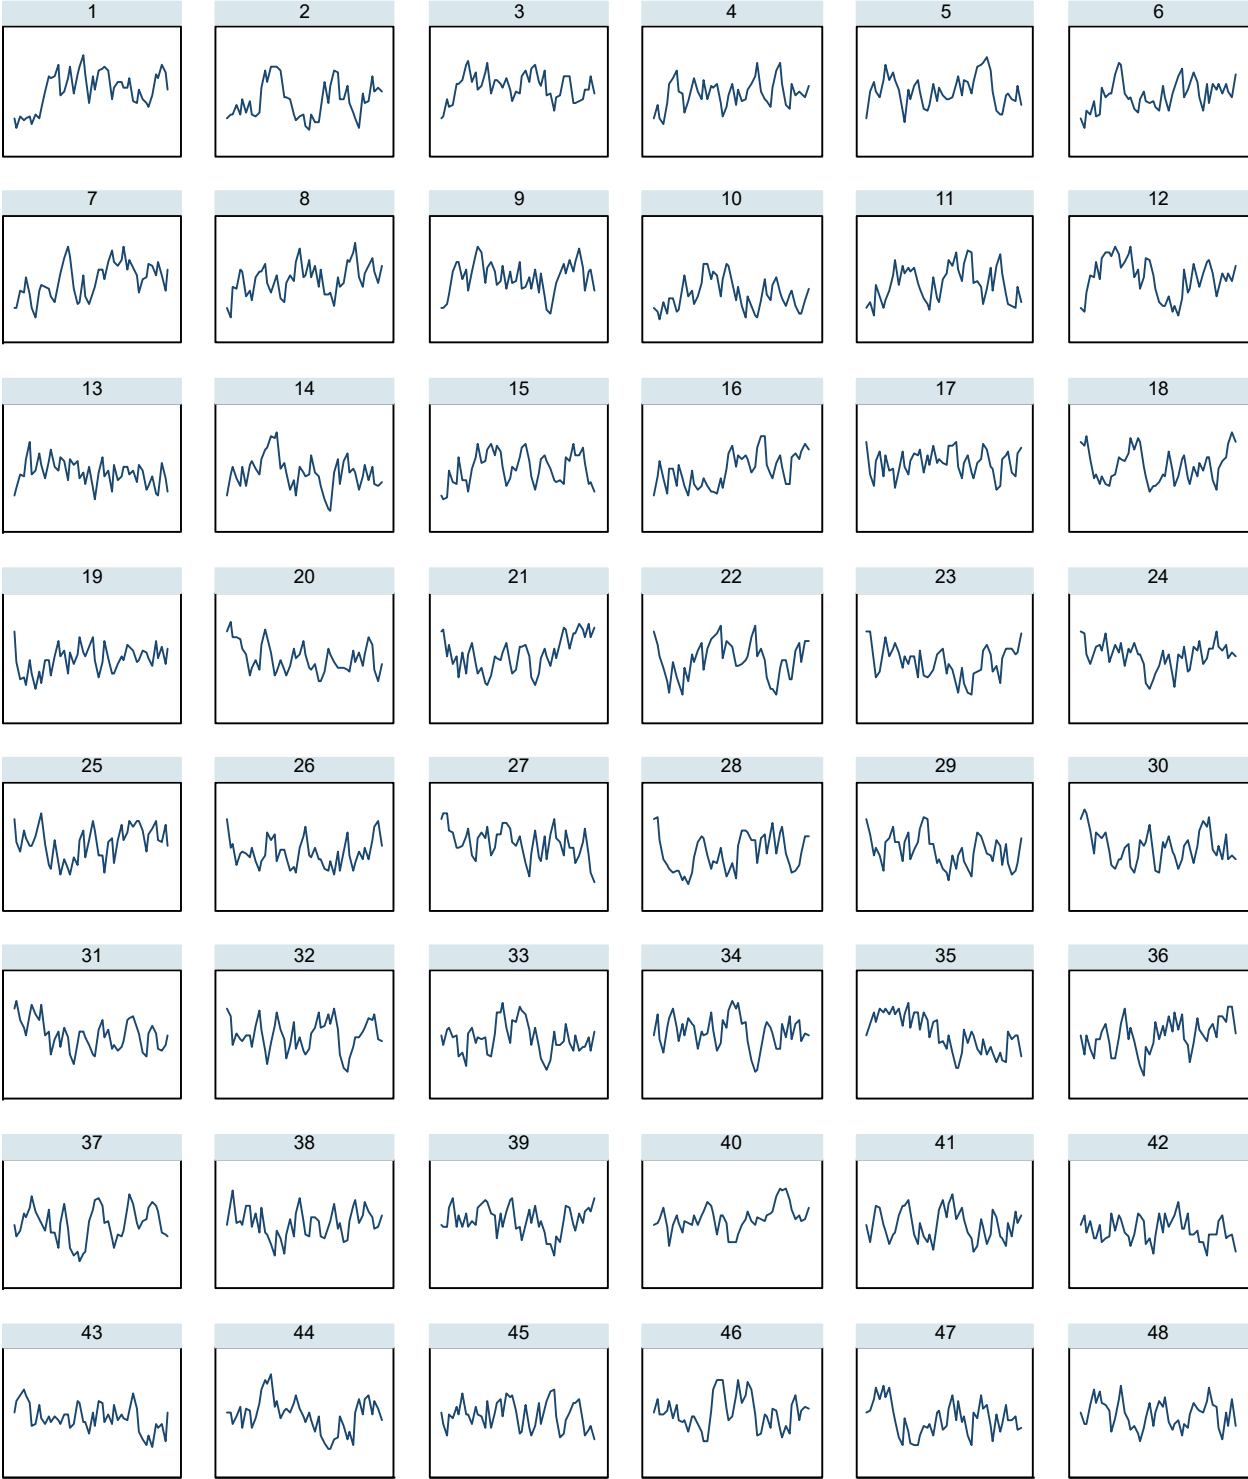


Figure A3: Prices in 48 markets.

B Behavior of Regret-Free Agent

The regret-free rational agent obtains utility $u(y_t)$ when she sells the asset at time t at the price y_t . In each period she estimates the expected future utility that takes into account her optimal choices and sells the asset if

$$u(y_t) \geq \max\{\mathbb{E}_{y_{t+1}}[u(y_{t+1})|y_t], \mathbb{E}_{y_{t+1}}[v_{t+2}|y_t]\}$$

where $v_{t+2} = \max\{\mathbb{E}_{y_{t+2}}[u(y_{t+2})|y_{t+1}], \mathbb{E}_{y_{t+2}}[v_{t+3}|y_{t+1}]\}$ and the value function in the last period T is $v_T = \mathbb{E}_{y_T}[u(y_T)|y_{T-1}]$.² In the experiment the price evolution is described by a Markov chain, thus, all expectations are conditional on the past price.

We show analytically that a risk-averse agent should optimally sell the asset at a lower price than a risk-neutral agent and risk-loving agent should sell at a higher price. Intuitively, an extremely risk-averse agent sells immediately at any price level as a sure outcome today outweighs an uncertain outcome tomorrow, whereas the certainty equivalent required by a risk-loving agent to sell at the same price is higher. Thus, we formulate a prediction concerning risk attitudes:

Proposition. *The optimal policy for an agent with CRRA preferences is to sell the asset above some threshold different for each period. Other things equal, the probability of selling the asset increases in the degree of risk aversion.*

Proof. An agent without regret sells if

$$u(y_t) \geq \max\{\mathbb{E}_{y_{t+1}}[u(y_{t+1})|y_t], \mathbb{E}_{y_{t+1}}[v_{t+2}|y_t]\} \quad (\text{B.1})$$

where $v_{t+2} = \max\{\mathbb{E}_{y_{t+2}}[u(y_{t+2})|y_{t+1}], \mathbb{E}_{y_{t+2}}[v_{t+3}|y_{t+1}]\}$ and $v_T = \mathbb{E}_{y_T}[u(y_T)|y_{T-1}]$. Assuming that agent has CRRA utility function, this implies that the selling rule is

$$y_t^{1-\rho} \geq \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|y_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|y_t]\}$$

where $\dot{v}_{t+2} = \max\{\mathbb{E}_{y_{t+2}}[y_{t+2}^{1-\rho}|y_{t+1}], \mathbb{E}_{y_{t+2}}[\dot{v}_{t+3}|y_{t+1}]\}$ and $\dot{v}_T = \mathbb{E}_{y_T}[y_T^{1-\rho}|y_{T-1}]$.

Let \tilde{v}_t denote the value function in inequality (B.1) with $u(y_t) = y_t$ and let \tilde{y}_t be the price at which a risk-neutral agent is indifferent whether to sell the asset or not:

$$\tilde{y}_t = \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\tilde{v}_{t+2}|\tilde{y}_t]\} \quad (\text{B.2})$$

Would a risk-seeking (averse) agent sell at the same value or continue? The answer depends on ρ . Agent sells at \tilde{y}_t if and only if

$$\tilde{y}_t^{1-\rho} \geq \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|\tilde{y}_t]\}. \quad (\text{B.3})$$

Plugging (B.2) into (B.3) we get

$$\max\{\mathbb{E}_{y_{t+1}}[y_{t+1}|\tilde{y}_t]^{1-\rho}, \mathbb{E}_{y_{t+1}}[\tilde{v}_{t+2}|\tilde{y}_t]^{1-\rho}\} \geq \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|\tilde{y}_t]\}.$$

This inequality holds (strictly) only for a risk-averse agent with $\rho \in (0, 1)$. To show this we start from period $T - 1$. Notice that

$$\begin{aligned} \mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho} &= \left(\sum_i \Pr\{y_{T-1,\mu}|y_{T-2}\} \mathbb{E}_{y_T}[y_T|y_{T-1,\mu}] \right)^{1-\rho} \\ &\quad \text{and} \\ \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}] &= \sum_i \Pr\{y_{T-1,\mu}|y_{T-2}\} \mathbb{E}_{y_T}[y_T^{1-\rho}|y_{T-1,\mu}] \end{aligned} \quad (\text{B.4})$$

²By design the participants in the last period are forced to sell at the current price.

where, given y_{T-2} , ι enumerates all possible values of y_{T-1} denoted by $y_{T-1,\iota}$. Next notice that the RHS's of (B.4) can be rewritten as

$$\left(\sum_{\iota} \Pr\{y_{T-1,\iota}|y_{T-2}\} \sum_{\xi_i} \Pr\{y_{T,\xi_i}|y_{T-1,\iota}\} y_{T,\xi_i} \right)^{1-\rho} = \left(\sum_{\zeta} p_{\zeta} y_{T,\zeta} \right)^{1-\rho} \quad (\text{B.5})$$

and

$$\sum_{\iota} \Pr\{y_{T-1,\iota}|y_{T-2}\} \sum_{\xi_i} \Pr\{y_{T,\xi_i}|y_{T-1,\iota}\} y_{T,\xi_i}^{1-\rho} = \sum_{\zeta} p_{\zeta} y_{T,\zeta}^{1-\rho}$$

respectively. Here ξ_i enumerates y_T for each ι and ζ enumerates all combinations of ι and ξ_i . Now, the RHS of the first equation in (B.5) is bigger than the RHS of the second by strict concavity of $(\cdot)^{1-\rho}$. Thus we can conclude that $\mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho} > \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]$ for all $\rho \in (0, 1)$.

Now we consider period $T - 2$. For some fixed y_{T-2} we want to show that

$$\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}|y_{T-2}]^{1-\rho}, \mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho}\} > \max\{\mathbb{E}_{y_{T-1}}[y_{T-1}^{1-\rho}|y_{T-2}], \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]\}. \quad (\text{B.6})$$

This is straightforward since we have just shown that $\mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho} > \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]$, which are the second terms of the max operators. According to the same strict concavity argument as above, $\mathbb{E}_{y_{T-1}}[y_{T-1}|y_{T-2}]^{1-\rho} > \mathbb{E}_{y_{T-1}}[y_{T-1}^{1-\rho}|y_{T-2}]$, the first terms of the max operators. Thus, LHS max operator has all terms bigger than corresponding terms of the RHS max operator, which proves that the inequality (B.6) holds.

Since (B.6) holds for all y_{T-2} , it is true that

$$\mathbb{E}_{y_{T-2}}[\tilde{v}_{T-1}|y_{T-3}]^{1-\rho} = \mathbb{E}_{y_{T-2}}[\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}|y_{T-2}]^{1-\rho}, \mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho}\}|y_{T-3}] > \mathbb{E}_{y_{T-2}}[\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}^{1-\rho}|y_{T-2}], \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]\}|y_{T-3}] = \mathbb{E}_{y_{T-2}}[\dot{v}_{T-1}|y_{T-3}].$$

This is a precursor to the one more step of the same derivation for period $T - 3$ as $\mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho} > \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]$ was for the period $T - 2$ step. Therefore, iterating this process, we show that (B.3) holds with strict inequality for all t as long as $\rho \in (0, 1)$. When the agent is risk-seeking, or $\rho < 0$, (B.3) holds strictly with the opposite sign. The proof is the same only with all $>$ replaced by $<$.

Next we show that for any admissible ρ and each period there is a unique threshold such that an agent with CRRA utility, who follows optimal policy, always sells above this threshold and always keep the asset below it. Notice that $\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|y_t] = \mathbb{E}_{\varepsilon}[(\alpha y_t + (1 - \alpha)\varepsilon)^{1-\rho}]$ is a strictly increasing continuous function of y_t .³ Consider $m(y_t) = \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|y_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|y_t]\}$. This is a function of y_t that for some y_t is equal to $\mathbb{E}_{\varepsilon}[(\alpha y_t + (1 - \alpha)\varepsilon)^{1-\rho}]$ and for some y_t to $\mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|y_t]$. Now, we can use the expressions $\dot{v}_{\tau} = \max\{\mathbb{E}_{y_{\tau}}[y_{\tau}^{1-\rho}|y_{\tau-1}], \mathbb{E}_{y_{\tau}}[\dot{v}_{\tau+1}|y_{\tau-1}]\}$ for all $\tau \geq t + 2$ to expand $\mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|y_t]$ into a sequence of expectations and max operators. Thus, eventually, $m(y_t)$ is a piecewise function that is equal to $\mathbb{E}_{\varepsilon}[(\alpha y_t + (1 - \alpha)\varepsilon)^{1-\rho}]$ or pieces of weighted averages of functions of the form

$$\mathbb{E}_{y_{t+1}}[\dots \mathbb{E}_{y_{\tau}}[y_{\tau}^{1-\rho}|y_{\tau-1}] \dots |y_t] = \mathbb{E}_{\varepsilon_{t+1}} \dots \mathbb{E}_{\varepsilon_{\tau}}[(\alpha^{\tau-t} y_t + (1 - \alpha^{\tau-t}) E_{\tau})^{1-\rho}] \quad (\text{B.7})$$

where E_{τ} is a weighted average of random variables $\varepsilon_{t+1}, \varepsilon_{t+2}, \dots, \varepsilon_{\tau}$. All functions in (B.7) are continuous and strictly increasing in y_t . Therefore, $m(y_t)$ is a continuous and strictly increasing since it is a series of max operators applied to weighted averages of continuous increasing functions. It is also true that m is strictly concave (convex) for $\rho \in (0, 1)$ ($\rho < 0$), which also follows from the fact that it is a series of max operators of weighted averages of strictly concave (convex) functions.

Now, we would like to know the relationship between $m(y_t)$ and $y_t^{1-\rho}$. This will tell us what the

³Here and below ε , possibly with sub-indexes, is a uniformly distributed random variable on $[0, 10]$.

optimal policy is. Notice that $m(0) > 0^{1-\rho}$ and $m(10) < 10^{1-\rho}$ since $m(y_t)$ consists of mean reverting expectations. So for low y_t the optimal policy is to keep the asset and for high y_t to sell. It is left to show that $m(y_t)$ crosses $y_t^{1-\rho}$ at a single point. Consider any point y where $y^{1-\rho} = m(y)$. We want to show that at this point the derivatives of $y^{1-\rho}$ and $m(y)$ are different. As was mentioned above, $m(y)$ is a weighted average of functions in (B.7). Thus,

$$y^{1-\rho} = \sum_t p_t \mathbb{E}_{\varepsilon_{t+1} \dots \varepsilon_{\tau_t}} [(\alpha^{\tau_t-t} y + (1 - \alpha^{\tau_t-t}) E_{\tau_t})^{1-\rho}] = \sum_t p_t \mathbb{E}_{\tau_t} [(\alpha^{\tau_t-t} y + (1 - \alpha^{\tau_t-t}) E_{\tau_t})^{1-\rho}] \quad (\text{B.8})$$

for some enumeration $\{p_t, \tau_t\}_t$ and with \mathbb{E}_{τ_t} being short for $\mathbb{E}_{\varepsilon_{t+1} \dots \varepsilon_{\tau_t}}$. Notice that the derivatives of functions (B.7) with respect to y_t are of the form $\alpha^{\tau-t}(1-\rho)\mathbb{E}_{\tau}(\alpha^{\tau-t}y_t + (1-\alpha^{\tau-t})E_{\tau})^{-\rho}$, since \mathbb{E}_{τ} transforms into a summation of the terms $(\alpha^{\tau-t}y + (1-\alpha^{\tau-t})E_{\tau})^{1-\rho}$ weighted with some probabilities and the derivative transcends the summation. Keeping this in mind let us rewrite (B.8) as

$$(1-\rho)y^{-\rho} = \sum_t p_t \alpha^{\tau_t-t} (1-\rho) \mathbb{E}_{\tau_t} [(\cdot)^{-\rho}] + \frac{1-\rho}{y} \sum_t p_t \mathbb{E}_{\tau_t} [(1-\alpha^{\tau_t-t}) E_{\tau_t} (\cdot)^{-\rho}]$$

where $(\cdot)^{-\rho}$ stands for $(\alpha^{\tau_t-t}y + (1-\alpha^{\tau_t-t})E_{\tau_t})^{1-\rho}$. This, in turn, can be seen in terms of derivatives

$$(1-\rho)y^{-\rho} = \frac{dm(y)}{dy} + \frac{1-\rho}{y} \sum_t p_t \mathbb{E}_{\tau_t} [(1-\alpha^{\tau_t-t}) E_{\tau_t} (\cdot)^{-\rho}].$$

Here LHS is the derivative of LHS of (B.8) at y and RHS is the derivative of m at y plus a positive number. Thus, at y the derivative of $y^{1-\rho}$ is higher than the derivative of $m(y_t)$. This implies that these two functions cross at a unique point: they cannot coincide on an interval, since then their derivatives would have been equal and they cannot cross on a disjoint set since this would have contradicted the strict concavity or convexity of m .

Thus, we have established that the optimal policy for any CRRA utility function is to sell above some unique threshold y_t and to keep the asset below it. Combining this observation with the result that risk the averse agent sells at a price where risk-neutral agent is indifferent and that the risk-seeking agent keeps the asset at that price, we can conclude that risk-averse agent must have the selling threshold at a price below the risk-neutral agent and the risk-seeking agent must have the threshold above it. Therefore, a risk-averse agent, given the same prices, sells before a risk-neutral agent and a risk-seeking agent sells after. ■

C Supplementary Graphs

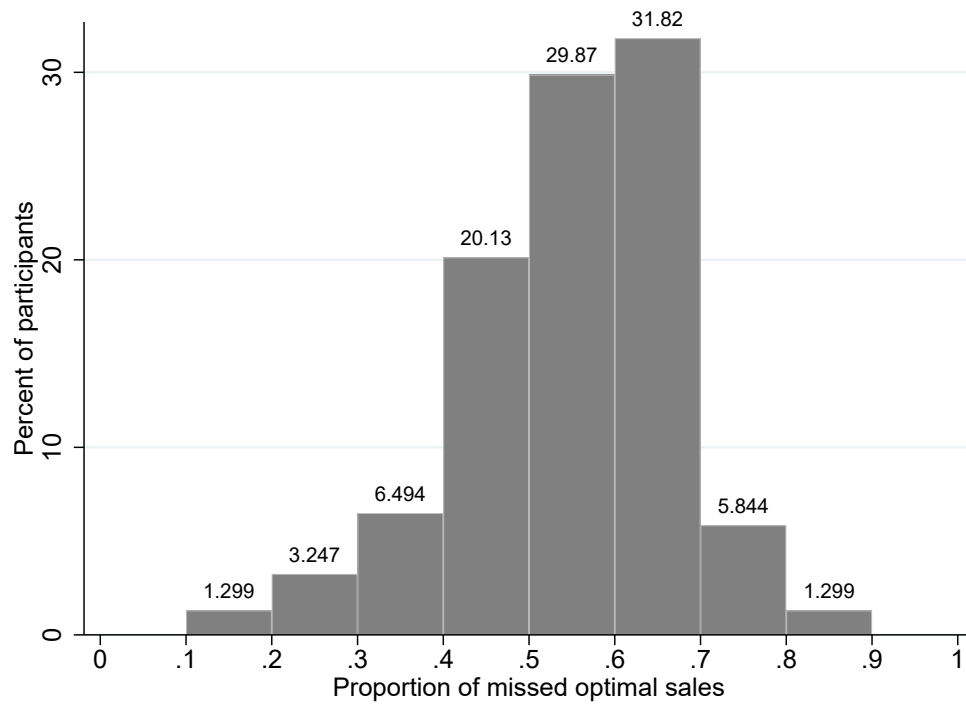


Figure C4: The distribution of participants by the proportion of missed optimal sales. The proportion is calculated separately for each participant.




Figure C5: The percentage of sales when the price reaches a new peak (dark grey) and when the price is below the current past peak (light grey). Only observations above the optimal selling price threshold of the risk-neutral no regret agent are considered. The error bars are $\pm 1SE$.




Figure C6: The ratio of the number of sales up to period t in the No Info to Info condition starting from period 33.

D Description of the Variables

Variable	Mean	Median	St. Dev.	Range	Definition
choice	0.94	1.00	0.24	{0, 1}	1 if the participant keeps the asset and 0 if she sells it
info				{0, 1}	1 if the market condition is Info and 0 if it is No Info
time	26.61	25.00	8.44	[16, 49]	Time period
price	4.79	4.86	1.43	[1.20, 8.36]	Current price
price ²	24.96	23.59	13.83	[1.43, 69.95]	Current price squared
future expected price	5.00	5.00	0.08	[3.48, 7.02]	Expected future price in period 50 conditional on the current price, the number of periods left, and the Markov process that generates the prices
past peak	7.58	7.52	0.59	[5.53, 8.56]	Highest price in the past
future expected peak	7.64	7.78	0.45	[3.48, 8.24]	Highest expected future peak conditional on the current price and time (see Appendix F for details)
hl	0.60	0.60	0.17	[0, 0.9]	Risk aversion parameter from Holt and Laury task (normalized from [0, 10] to [0, 1]). 1 is very risk-averse, 0 is very risk-seeking
early				{0, 1}	1 for first 25, 28, 30, or 32 markets depending on specification, 0 otherwise

Table D1: Variables used in the regressions in Tables 1, E2, and E3 (Appendix E). The statistics refer to all periods when a choice is made (periods 16 to 49).

E Additional Regressions

Pr[choice = keep]	I	II	III	IV	V	VI
price	-0.497*** (0.146)	-0.416*** (0.161)	-0.237 (0.145)	-0.243* (0.145)	-0.243* (0.145)	-0.243* (0.145)
price ²	-0.102*** (0.013)	-0.113*** (0.014)	-0.136*** (0.013)	-0.136*** (0.013)	-0.136*** (0.013)	-0.136*** (0.013)
time	-0.088*** (0.004)	-0.090*** (0.004)	-0.082*** (0.005)	-0.083*** (0.005)	-0.083*** (0.005)	-0.083*** (0.005)
future expected price	1.423*** (0.230)	1.223*** (0.251)	1.226*** (0.205)	1.218*** (0.203)	1.218*** (0.203)	1.217*** (0.203)
past peak			0.522*** (0.041)	0.617*** (0.053)	0.617*** (0.053)	0.643*** (0.140)
future expected peak			0.343*** (0.080)	0.272*** (0.093)	0.271*** (0.093)	0.247 (0.244)
past peak × info				-0.208*** (0.072)	-0.208*** (0.072)	-0.209*** (0.072)
future expected peak × info				0.132* (0.074)	0.133* (0.074)	0.133* (0.074)
info				0.511 (0.754)	0.552 (0.757)	0.553 (0.756)
hl		-0.649** (0.301)	-0.646** (0.304)	-0.651** (0.305)	-0.620** (0.307)	-0.622 (2.902)
info × hl					-0.062 (0.205)	-0.060 (0.206)
hl × past peak						-0.038 (0.201)
hl × future expected peak						0.038 (0.353)
constant	4.525*** (1.161)	5.939*** (1.274)	-1.044 (1.090)	-1.116 (1.214)	-1.135 (1.222)	-1.139 (2.259)
<i>N</i>	112,137	89,951	89,951	89,951	89,951	89,951

Table E2: Random effects logit regression of the choice to keep the asset with risk preferences. choice is zero at the time the participant sells the asset and one otherwise. Observations are all periods in all markets for all participants in which they made a choice (periods 16 to 49). Participants whose choices in Holt-Laury task were inconsistent with expected utility maximization were dropped. Errors are clustered by participant and robust.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

	I	II	III	IV
Pr[choice = keep]	early: 25	early: 28	early: 30	early: 32
price	-0.331** (0.134)	-0.323** (0.135)	-0.280** (0.133)	-0.292** (0.134)
price ²	-0.124*** (0.012)	-0.124*** (0.012)	-0.128*** (0.012)	-0.128*** (0.012)
time	-0.092*** (0.005)	-0.120*** (0.006)	-0.130*** (0.006)	-0.112*** (0.006)
future expected price	1.470*** (0.191)	1.646*** (0.191)	1.621*** (0.190)	1.484*** (0.186)
past peak	0.602*** (0.045)	0.588*** (0.045)	0.604*** (0.045)	0.586*** (0.045)
future expected peak	0.190** (0.084)	0.057 (0.088)	0.022 (0.096)	0.179** (0.088)
past peak × info	-0.204*** (0.065)	-0.196*** (0.065)	-0.212*** (0.065)	-0.205*** (0.065)
future expected peak × info	0.089 (0.078)	0.095 (0.079)	0.227** (0.088)	0.168** (0.085)
info	0.728 (0.726)	0.624 (0.714)	-0.136 (0.765)	0.199 (0.755)
info × early	0.176** (0.072)	0.159** (0.081)	-0.065 (0.087)	0.022 (0.089)
early	-0.239*** (0.070)	-0.718*** (0.072)	-0.854*** (0.078)	-0.657*** (0.074)
constant	-1.670 (1.086)	-0.385 (1.095)	0.185 (1.121)	-0.704 (1.090)
<i>N</i>	112,137	112,137	112,137	112,137

Table E3: The logit regressions support the intuition in Figure 4B. The dummy variable early is 1 if the current period is smaller or equal than the value specified in each column title and 0 otherwise. Participants in the Info condition sell less often early on because of the possibility of future regret: the coefficient on the interaction of info and early is significant and positive until Column III. Observations are all periods in all markets for all participants in which they made a choice (periods 16 to 49). Errors are clustered by participant and robust. The descriptions of all the variables can be found in Appendix D.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

F The Computation of Future Regret

At period t future regret is defined as the expectation of the highest order statistic of the future $T - t$ prices. At every period $t \in \{2, \dots, T\}$, $y_{t+1} = \alpha y_t + (1 - \alpha)u_t$ is observed, where u_t is an i.i.d. random draw from the uniform distribution on $[a, b]$. We use the notation y_t^k to indicate the price expected in period t given the current price in period k . T , α and y^k (price at time k) are known. Assume a given period $k \in \{1, \dots, T - 2\}$, noting that the expected future peak in the period before the last is just the expectation of the price in the next period. Then we can recover the expected price for any future period beyond the current period ($\forall t > k$) with the following formula:

$$y_t^k = \alpha^{t-k} y_k + (1 - \alpha) \sum_{j=1}^{t-k} u_{t-j} \alpha^{j-1} \quad (\text{F.1})$$

The distribution of y_t^k is

$$\begin{aligned} P\{y_t^k \leq v\} &= P\{\alpha^{t-k} y_k + (1 - \alpha) \sum_{j=k}^{t-1} u_{t-j} \alpha^{j-1} \leq v\} \\ &= F_{(t-k)}(v) = \int_0^v f_{(t-k)}(s) ds \end{aligned}$$

where $f_{(t-k)}(s)$ is the pdf of the sum of $(t - k)$ uniform distributions with different supports. The support of this distribution is $(\alpha^{t-k} y_k, \alpha^{t-k} y_k + 10(1 - \alpha) \sum_{j=1}^{t-k} \alpha^{j-1})$. This is again when all u 's are 0 or all u 's are 10. Note that when $t - k = 1$ $f_{(1)}(s) = \frac{1}{\alpha y_k + (1 - \alpha)10 - \alpha y_k} = \frac{1}{(1 - \alpha)10}$ and $F_{(1)}(s) = \frac{s - \alpha y_k}{(1 - \alpha)10}$.

The expected future peak is computed as:

$$\begin{aligned} \text{Future peak}_{\text{period } k} &= \int_0^{10} v d \prod_{j=1}^{T-k} F_{(j)}(v) \\ &= \int_0^{10} v \sum_{j=1}^{T-k} f_{(j)}(v) \prod_{h \neq j}^{T-k} F_{(h)}(v) dv \\ &= \int_0^{10} v \sum_{j=1}^{T-k} f_{(j)}(v) \prod_{h \neq j}^{T-k} \int_0^v f_{(h)}(s) ds dv \end{aligned}$$

To derive $f_{(t-k)}(v)$ analytically we use recent results in the statistical literature ([Potuschak and Muller, 2009](#)). For simplicity assume that $k = 1$. In fact, the random variable in (F.1) is the sum of independent uniformly distributed $[0, 10]$ random variables times $(1 - \alpha) \times \alpha^{j-1}$, plus $\alpha^{t-1} \frac{y_1}{t-1}$, which is equal to the summation of $t - 1$ uniformly distributed random variables in $[\alpha^{t-1} \frac{y_1}{t-1}, \alpha^{t-1} \frac{y_1}{t-1} + 10(1 - \alpha) \alpha^{j-1}]$, $\forall j \in \{1, \dots, t - 1\}$. According to Potuschak and Muller (2009, section 2.2.2, page 180), the density is

$$f_{(n)}(s) = \frac{1}{2^n (n - 1)! \prod_k a_k} \sum_{j=1}^{2^n} \sigma_j \max\{\underline{a} \cdot \underline{\varepsilon}_j - |s - \sum_k c_k|, 0\}^{n-1} \quad (\text{F.2})$$

where \cdot indicates the dot product, lower bar means vector, $a = \{5(1 - \alpha), 5\alpha(1 - \alpha), 5\alpha^2(1 - \alpha), \dots, 5\alpha^{t-1}(1 - \alpha)\}$, $c = \{\alpha^{t-1} \frac{y_1}{t-1} + 5(1 - \alpha), \alpha^{t-1} \frac{y_1}{t-1} + 5\alpha(1 - \alpha), \alpha^{t-1} \frac{y_1}{t-1} + 5\alpha^2(1 - \alpha), \dots, \alpha^{t-1} \frac{y_1}{t-1} + 5\alpha^{t-1}(1 - \alpha)\}$, $\forall 1 \leq j \leq t - 1$. σ_j and $\underline{\varepsilon}_j$ are matrices which deal with positive and negative signs (see [Potuschak and Muller \(2009\)](#)). We can rewrite the distribution as follows:

$$P\{y_t^1 \leq v\} = F_{(t)}(v) = \int_0^v f_{(t-1)}(s) ds$$

The support of this distribution is $[\alpha^{t-1}y_1, \alpha^{t-1}y_1 + 10(1-\alpha)\sum_{j=1}^{t-1}\alpha^{j-1}]$. Note that $f_{(1)}(s) = \frac{1}{\alpha y_1 + (1-\alpha)10 - \alpha y_1} = \frac{1}{(1-\alpha)10}$ and $F_{(1)}(s) = \frac{s - \alpha y_1}{(1-\alpha)10}$.

F.1 Normal Approximation

(F.2) is problematic, because, as the number of uniform RVs to be summed increases, the denominator goes to zero since $a_k \rightarrow 0$. This makes estimation intractable. Another unappealing feature of this equation is that computation is extremely slow. Therefore, we follow [Potuschak and Muller \(2009\)](#) who proposed to approximate $f_{(n)}(v) = f_{(t-k)}(v)$ with the following normal distribution:

$$y_t^k \sim \mathcal{N}\left(\sum_k c_k, \sum_k \frac{(2 \cdot a_k)^2}{12}\right)$$

The approximation is based on the fact that the sum of uniform distributions is centered around $\sum_k c_k$ with variance $\frac{1}{12}(b-a)^2$, where b and a are the upper and lower bounds of the support of the sum of uniform distributions.

It can be shown that the sum of such i.n.d. uniformly distributed random variables converges to a normal distribution by the Liapounov Central Limit Theorem. The condition for convergence is:

$$\lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N E[|y_i - \mu_i|^{2+\beta}]}{(\sum_{i=1}^N \sigma_i^2)^{\frac{2+\beta}{2}}} = 0,$$

for some choice of $\beta > 0$, where $E[y_i] = \mu_i$ and $V[X_i] = \sigma_i^2$. To see this assume $\beta = 1$ for simplicity and denote $X_i = y_i - \mu_i$. Because $\mu_i = c_i$ and the support of y_i is $[c_i - a_i, c_i + a_i]$, X_i is uniformly distributed in the interval $[-a_i, a_i] = [-5(1-\alpha)\alpha^{i-1}, 5(1-\alpha)\alpha^{i-1}]$. The numerator of the CLT condition involves $E[|X_i|^3] = \int_{-a_i}^{a_i} |s|^3 f_i(s) ds = \int_{-a_i}^{a_i} |s|^3 \frac{1}{2a_i} ds$. Solving the integral we get:

$$\begin{aligned} E[|X_i|^3] &= \frac{1}{2a_i} \frac{1}{4a_i} s^4 \text{sgn}(s) \Big|_{-a_i}^{a_i} \\ &= \frac{125}{4} (1-\alpha)^3 \alpha^{3(i-1)} \end{aligned}$$

Therefore, the numerator is $\frac{125}{4} \sum_i^N (1-\alpha)^3 \alpha^{3(i-1)}$. Similarly, the denominator can be rewritten using the formula for the variance of the normal distribution as $(\frac{25}{3})^{\frac{3}{2}} (\sum_i^N (1-\alpha)^2 \alpha^{2(i-1)})^{\frac{3}{2}}$ (use the fact that $\sigma_i^2 = \frac{1}{12}(c_i + a_i - (c_i - a_i))^2 = \frac{1}{12}(2 \times a_i)^2$). Taking the ratio of these two quantities, the result is $W \times \frac{\sum_i^N (1-\alpha)^3 \alpha^{3(i-1)}}{(\sum_i^N (1-\alpha)^2 \alpha^{2(i-1)})^{\frac{3}{2}}}$, where $0 < W < 1$ is a constant. Finally, we can establish that:

$$\begin{aligned} \lim_{N \rightarrow \infty} &= \frac{\sum_{i=1}^N E[|X_i|^3]}{(\sum_{i=1}^N \sigma_i^2)^{\frac{3}{2}}} \\ &= W \times \frac{\sum_{i=1}^N (1-\alpha)^3 \alpha^{3(i-1)}}{(\sum_{i=1}^N (1-\alpha)^2 \alpha^{2(i-1)})^{\frac{3}{2}}} \\ &= 0 \end{aligned}$$

because the denominator contains positive interaction terms. Therefore, $\sum y_i \sim \mathcal{N}(\sum_k c_k, \sum_k \frac{(2 \times a_k)^2}{12})$.




Figure F7: pdf, sum of 3 uniform RVs




Figure F8: pdf, sum of 13 uniform RVs




Figure F9: CDF, sum of 3 uniform RVs




Figure F10: CDF, sum of 13 uniform RVs

G Discretization of the State Space and Transition Matrix

After discretization of the state space, the process describing the evolution of the price at each period of time can be represented by a discrete Markov chain. In fact, the only determinant of price in the next period is the price in the previous period. The discretization is done following [Tauchen \(1986\)](#). See also [Aguirregabiria and Magesan \(2016, page 23\)](#).

The stochastic shock follows the following AR(1) process:

$$y_{i,t+1} = \mu + \rho y_{i,t} + \varepsilon \quad (\text{G.1})$$

where $y_{i,t+1}, y_{i,t}$ are the prices for participant $i = \{1, \dots, N\}$ at time $t + 1$ and t respectively, and $\varepsilon \sim N(0, \sigma^2)$. This panel structure is composed of 48 sequences (the individual dimension) and 50 periods (the time dimension). $\hat{\mu}$ and $\hat{\rho}$ are found using the covariance estimator. The estimates are $\hat{\mu} = 1.97$, $\hat{\rho} = 0.60$ and $\hat{\sigma} = 1.16$. The estimate of ρ is very close to the parameter α which updates the price from period y_t to y_{t+1} ($\alpha = 0.6$).

Let $\{y^1, \dots, y^K\}$ denote the support of the discretized variable $\tilde{Y}_{i,t}$, where $y^1 > y^2 > \dots > y^{K-1} > y^K$ with $K = 400$ are the points in the support. [Tauchen \(1986\)](#) suggests using

$$y^K = \frac{\mu}{1-\rho} + m \times \left(\frac{\sigma^2}{1-\rho^2} \right)^{\frac{1}{2}}$$

$$y^1 = \frac{\mu}{1-\rho} - m \times \left(\frac{\sigma^2}{1-\rho^2} \right)^{\frac{1}{2}}$$

and y^k are $K - 2$ equidistant points within y^K and y^1 , such that the distance between any two points is ω . m is the density of the K points (m is set to 3). This choice of the parameters results in a support with lower bound (y^1) equal to €0.59 ca., upper bound (y^{200}) equal to €9.32 ca., and interval between adjacent points (ω) equal to €0.02 ca.

The probability of transitioning from state y to y' is defined as $p_{i,j} = Pr(y' = y^j | y = y^i)$, which describes the element in the transition matrix in row i and column j . Because of the normality assumption,⁴ the transition probability to a state k , $1 < k < K$, from i is:

$$p_{i,k} = \Phi\left(\frac{y^k + \frac{\omega}{2} - \hat{\mu} - \hat{\rho}y^i}{\hat{\sigma}}\right) - \Phi\left(\frac{y^k - \frac{\omega}{2} - \hat{\mu} - \hat{\rho}y^i}{\hat{\sigma}}\right)$$

which can be thought as the probability that $\rho y^i + \varepsilon \in [\rho y^j - \frac{\omega}{2}, \rho y^j + \frac{\omega}{2}]$. Analogously, the transition probability to the first and last state are:

$$p_{i,1} = \Phi\left(\frac{y^1 + \frac{\omega}{2} - \hat{\mu} - \hat{\rho}y^i}{\hat{\sigma}}\right)$$

$$p_{i,K} = 1 - \Phi\left(\frac{y^K - \frac{\omega}{2} - \hat{\mu} - \hat{\rho}y^i}{\hat{\sigma}}\right)$$

[Tauchen \(1986\)](#) shows that this conditional distribution converges in probability to the true conditional distribution for the stochastic process in (G.1). In fact, it can be shown that such a discretization implies a stationary distribution with AR(1) parameters of $\rho = 0.60$ (equal to the α used in the experiment).

⁴The standardization implies that the distribution is a standard normal.

H Full Derivation of the Dynamic Discrete Choice Model

In this section we present the dynamic discrete choice model that will be used for the structural estimation of the risk and regret parameters of the utility function. The following derivations are also sketched in Section 5. Analogously to the logit panel regressions in Table 1, participants' choice between selling the asset or continuing still follows a threshold rule. However, they now take into account the Markovian nature of the problem. In particular, a participant's intertemporal utility is

$$\mathbb{E} \left\{ \sum_{t=1} \beta^{t-1} u_t^d(x_t) + \varepsilon_t^{d_t} \right\}$$

where $\beta \in (0,1)$ is a discount factor and $\varepsilon_t^{d_t}$ is an error term. As is customary in the dynamic discrete choice literature (Abbring, 2010; Aguirregabiria and Mira, 2010) it is assumed to be known and equal for all participants.⁵ d is the participant's binary choice at time $t \leq T$:

$$d_t = \begin{cases} 1, & \text{keep the asset} \\ 0, & \text{sell the asset.} \end{cases}$$

$u^{d_t}(x_t)$ is the payoff after choosing alternative d_t ; the observables are described by the realization of x_t , which is a tuple consisting of the current price y_t , the past maximum $s_{p,t}$, and the expected future maximum price $s_{f,t}$. We use a utility function which incorporates past and future regret as well as risk preferences. That is, we are interested in a utility function of the type $u(x_t) = U(x_t) - R(x_t)$, where $U(x_t)$ is a consumption utility function and $R(x_t)$ measures regret.

The flow (per period) payoff from choice d_t at period t is $u^{d_t} + \varepsilon_t^{d_t}$ where the error term $\varepsilon_t^{d_t}$ is independent of x_t . As is customary, the error term is assumed to be $\varepsilon_t^{d_t} = \tilde{\varepsilon}^{d_t} - \sigma_\varepsilon \gamma$ where $\tilde{\varepsilon}^{d_t}$ is distributed Type I extreme value with location parameter equal to zero and scale parameter $\sigma_\varepsilon = 1$.⁶ By the properties of the Type I extreme value distribution, the mean of $\tilde{\varepsilon}^{d_t}$ is γ (the Euler's constant). $\varepsilon_t^{d_t}$ is therefore mean zero. Given these preliminaries, denote by $V_t(x_t, \varepsilon_t) = \max_{d_t \in \{0,1\}} \{v_t^{d_t}(x_t) + \varepsilon_t^{d_t}\}$ the value function at the beginning of period t with $\varepsilon_t^{d_t} = \{\varepsilon_t^0, \varepsilon_t^1\}$ and define the *alternative specific value function* (ASVF) for option $d_t \in \{0,1\}$ at time t as:

$$v_t^{d_t}(x_t) = \begin{cases} 0 + \beta \mathbb{E}\{v_{t+1}(x_{t+1}) | x_t, d_t = 1\} & \text{if } d_t = 1 \text{ (keep)} \\ u^0(x_t) & \text{if } d_t = 0 \text{ (sell)} \end{cases} \quad (\text{H.1})$$

where the payoff of continuing is normalized to 0. Note that choosing to sell the asset implies null future payoffs (*terminating action*). The ex-ante value function in (H.1), can be rewritten as the expectation over the error term, ε_t , of the value function at time t

$$v_t(x_t) \equiv \int V_t(x_t, \varepsilon_t) d\Lambda(\varepsilon_t)$$

where $\Lambda(\cdot)$ is the logit distribution and $V_t(x_t, \varepsilon_t) = \max_{d_t \in \{0,1\}} \{v_t^{d_t}(x_t) + \varepsilon_t^{d_t}\}$. Define the *alternative specific value function* (ASVF) as:

$$v_t^{d_t}(x_t) = u^{d_t}(x_t) + \beta \mathbb{E}\{v_{t+1}(x_{t+1}) | x_t, \}, \quad d_t \in \{0,1\}. \quad (\text{H.2})$$

⁵Identification of the discount factor is possible only under an exclusion restriction (Magnac and Thesmar, 2002), and its estimation is generally hard. In order to circumvent this issue, we show that the estimations are robust to different values of β .

⁶The standard deviation of the error term is not identifiable in general, and therefore assumed to be equal to 1.

Because of the property of the Bellman equation, the optimal decision rule can be summarized as follows:

$$d_t = \begin{cases} 1 & \text{if } v_t^1(x_t) - v_t^0(x_t) \geq \varepsilon_t^0 - \varepsilon_t^1 \text{ at } t \\ 0 & \text{otherwise} \end{cases}$$

where $v^{d_t}(\cdot)$ is defined as in (H.2). Denote the *Conditional Choice Probability (CCP)* of selling (action 0) in period t as $\Pr\{d_t = 0|x_t\} \equiv p_t^0(x_t)$:

$$p_t^0(x_t) = \frac{\exp(v_t^0(x_t))}{\exp(v_t^0(x_t)) + \exp(v_t^1(x_t))} = \frac{1}{1 + \exp(v_t^1(x_t) - v_t^0(x_t))}. \quad (\text{H.3})$$

Therefore $p_t^0(x_t) = \Lambda\{v_t^1(x_t) - v_t^0(x_t)\}$. Due to the properties of the logit distribution $\Lambda\{\cdot\}$:

$$\phi(p_t^0(x_t)) \equiv \ln(1 - p_t^0(x_t)) - \ln(p_t^0(x_t)) \equiv v_t^1(x_t) - v_t^0(x_t). \quad (\text{H.4})$$

$\phi(\cdot)$ is estimable from choice data using (H.3) and (H.4) for each period t . Hence the difference in the alternative specific value functions, $v_t^1(x_t) - v_t^0(x_t)$, is known for every t . We can write the two ASVFs as follows:

$$\begin{aligned} v_t^0(x_t) &= u^0(x_t) \\ v_t^1(x_t) &= 0 + \beta \int_{\mathcal{X}} \int_{\varepsilon} \max\{v_{t+1}^0(x_{t+1}) + \varepsilon_{t+1}^0, v_{t+1}^1(x_{t+1}) + \varepsilon_{t+1}^1\} d\Lambda(\varepsilon) dF(x_{t+1}|x_t), \end{aligned} \quad (\text{H.5})$$

where the expectation in the second equation is only over the continuation alternative (1), because the transition matrix in case the absorbing choice (0) is chosen is zero for all x_t (i.e. $F(x_{t+1}|x_t, d_t = 0) = 0$). The estimation is based on the difference of the two ASVFs in (H.5):

$$v_t^1(x_t) - v_t^0(x_t) = -u(x_t) + \beta \int_{\mathcal{X}} \int_{\varepsilon} \max\{v_{t+1}^0(x_{t+1}) + \varepsilon_{t+1}^0, v_{t+1}^1(x_{t+1}) + \varepsilon_{t+1}^1\} d\Lambda(\varepsilon) dF(x_{t+1}|x_t) \quad (\text{H.6})$$

Notice that the LHS of (H.6) can be computed directly from the data using (H.4). The properties of the logit distribution are helpful to rewrite equation H.6 in a form that allows for estimation by non-linear least squares. In fact, the ASVF for continuing (second equation in H.5) can be rewritten as follows

$$\begin{aligned} v_t^1(x_t) &= \beta \int_{\mathcal{X}} \int_{\varepsilon} \max\{v_{t+1}^0(x_{t+1}) + \varepsilon_{t+1}^0, v_{t+1}^1(x_{t+1}) + \varepsilon_{t+1}^1\} d\Lambda(\varepsilon) dF(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}} \gamma + \log(\exp(v_{t+1}^0(x_{t+1}) - \gamma) + \exp(v_{t+1}^1(x_{t+1}) - \gamma)) dF(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}} \gamma + \log\left(\left(1 + \exp(v_{t+1}^1(x_{t+1}) - v_{t+1}^0(x_{t+1}))\right) \exp(v_{t+1}^0(x_{t+1}) - \gamma)\right) dF(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}} \left(u^0(x_{t+1}) - \log(p_{t+1}^0(x_{t+1}))\right) dF(x_{t+1}|x_t) \end{aligned}$$

where d_{t+1} is the decision in the next period and γ is the Euler's constant. The last row uses (H.3). Therefore the difference of the two ASVFs in (H.6) becomes

$$v_t^1(x_t) - v_t^0(x_t) = -u^0(x_t) + \beta \int_{\mathcal{X}} \left(v_{t+1}^0(x_{t+1}) - \log(p_{t+1}^0(x_{t+1}))\right) dF(x_{t+1}|x_t).$$

By replacing the dependent variable in the last equation with $\phi(p_t^0(x_t))$ and by discretizing the state space

\mathcal{X} the objective function becomes

$$\phi(p_t^0(x_t)) = -u^0(x_t) + \beta \sum_{x_{t+1} \in \mathcal{X}} \left(v_{t+1}^0(x_{t+1}) - \log(p_{t+1}^0(x_{t+1})) \right) f(x_{t+1}|x_t)$$

which concludes the derivation.

Note that the regret components are functions of price (y_t is the only random variable) and time. In fact, $s_{p,t} = \max_{\tau \leq t} y_\tau$ and $s_{f,t} = g(y_t, t)$, where g is a known function that is increasing in the first argument and decreasing in the second.⁷ Therefore,

$$Pr\{y_{t+1}, s_{p,t+1}, s_{f,t+1} | x_t, d_t = 1\} = f(y_{t+1}, s_{p,t+1}, s_{f,t+1} | y_t, s_{p,t}, s_{f,t}) = f(y_{t+1}, s_{p,t+1}, s_{f,t+1} | y_t, s_{p,t}).$$

The transition of the past peak is fully defined by the future price: if $y_{t+1} \geq s_{p,t}$ then $s_{p,t+1} = y_{t+1}$ and $s_{p,t+1} = s_{p,t}$ otherwise. For clarity, consider the following example: given the information available at period $t < T$, the expected utility from keeping the asset one period longer, in the Info condition, is given by

$$\mathbb{E}[u(x_{t+1}) | x_t] = \sum_{y_{t+1}} [\mathbb{1}_{\{y_{t+1} \geq s_{p,t}\}} u(y_{t+1}, y_{t+1}, g(\cdot)) + \mathbb{1}_{\{y_{t+1} < s_{p,t}\}} u(y_{t+1}, \max_{\tau \leq t} y_\tau, g(\cdot))] f(y_{t+1} | y_t).$$

Finally, the transition of the expected future peak is completely determined by the price and time according to the function $g(y_t, t)$.

⁷ $g(\cdot)$ is not strictly monotonic in the two arguments because of the discretization imposed to the data.

I Additional Estimations of the Structural Model

In this section, we propose a series of robustness checks employing different specifications of the utility function compared to that in the main text (see Section 4) to study how ω , α , and λ affect decisions.

I.1 No Future Regret in No Info Condition

Here we estimate several models where we assume that participants do not experience future regret in the No Info condition. Suppose that the per-period utility function

$$u(y_t, s_{p,t}, s_{f,t}) = \pi U(y_t) - R(s_{p,t}, s_{f,t}) \quad (\text{I.1})$$

includes a regret function $R(\cdot, \cdot)$ defined by

$$R(s_{p,t}, s_{f,t}) = \mathbb{1}_{\{\text{No Info}\}} \omega_{NI} s_{p,t} + \mathbb{1}_{\{\text{Info}\}} (\omega_I s_{p,t} + \alpha_I s_{f,t}). \quad (\text{I.2})$$

The arguments of the regret function are the past and expected future peaks and the market conditions are denoted by the subscripts “NI” for No Info and “I” for Info. The indicator function distinguishes the utility derived in one condition from the other. If ω_{NI} , ω_I , and α_I are not significantly different from zero, the participants are categorized as regret neutral.

The results from estimating (5.3) are shown in the first three columns of Table I4 and are obtained by nonlinear least squares on the dataset including periods $t \in \{16, \dots, 48\}$.⁸ The results are robust across different designs and discount factors.

	Estimation of (I.2)			Estimation of (I.3)		
	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
$\hat{\pi}$	1.892*** (0.011)	1.890*** (0.011)	1.888*** (0.011)	1.884*** (0.011)	1.882*** (0.011)	1.880*** (0.011)
$\hat{\omega}_{NI}$	0.313* (0.188)	0.379* (0.185)	0.464* (0.180)	0.364* (0.188)	0.428** (0.185)	0.508*** (0.181)
$\hat{\omega}_I$	-0.073 (0.195)	0.033 (0.192)	0.145 (0.188)	1.539*** (0.400)	1.636*** (0.359)	1.712*** (0.326)
$\hat{\alpha}_I$	0.174** (0.076)	0.221** (0.075)	0.262*** (0.074)	1.488*** (0.295)	1.552*** (0.262)	1.595*** (0.238)
$\hat{\lambda}_I$				-0.206*** (0.045)	-0.215*** (0.041)	-0.221*** (0.038)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table I4: The estimation of (5.3) with the regret terms as in (I.2) and (I.3) in periods 16 to 48 for different values of the discount factor β . The utility is assumed linear. Standard errors are in parenthesis.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

The estimation of (I.2) in Table I4 shows that participants are both past and future regret averse. In particular, in the specification (I.2) past regret is significant in the No Info condition, while future regret

⁸For consistency period 49 is dropped because choices taken in this period are directly affected by the fact that participants are *forced* to sell in period 50. This marginally shrinks the dataset from 112,137 to 111,613 observations. Including period 49 does not change the results. Note that the CCP must still be computed for period 49.

is significant in the Info condition which means that our participants are also influenced by future regret avoidance. Notice that past regret is not significant in the estimation of (I.2). The absence of the effect of the past peak is surprising given the discussion in Section 3 and the regression analysis in Tables 1, E2, and E3, which shows the centrality of the $s_{p,t}$ term for both conditions. The reason for this might be that the model is missing an important interaction between the past and the future regret: they might reinforce or inhibit each other. Such an interaction was previously exposed in the discussion of Figure 4B and regression E2 in Section 3 when we compared the decisions to sell early and late in the two conditions. Its presence was also used to explain changes in the probability of selling the asset with different values of the past and future peak in Figure 6 in Section 6.1. To our knowledge, no one in the literature pointed out the importance of this interaction. To account for it we reformulate the regret function to include an interaction term in the Info condition as in (4.1):

$$R(s_{p,t}, s_{f,t}) = \mathbb{1}_{\{\text{No Info}\}} \omega_{NI} s_{p,t} + \mathbb{1}_{\{\text{Info}\}} (\omega_I s_{p,t} + \alpha_I s_{f,t} + \lambda_I s_{p,t} s_{f,t}). \quad (\text{I.3})$$

The interaction term captures the cross-partial derivative of the regret function, which allows us to understand the degree of complementarity or substitutability of the two peaks. The last three columns of Table I4 display the results of the estimation of (5.3) with the regret term (I.3). The results confirm that participants are averse to past regret in both conditions and to future regret in the Info condition. In addition, it indicates a pattern of substitution between the two reference points, as $\hat{\lambda}_I < 0$. Notice also that the estimate of the coefficient on the consumption utility is higher than either the coefficient on past or future regret in the Info condition. As we will see below this implies that participants care more about consumption than regret.⁹

I.2 Regret and Risk Preferences

This section reports estimates for several models displaying different parameterization of the regret-averse utility function and risk-aversion. All estimations are consistent with the findings displayed in Section 6. The tables below show NLS estimates assuming the following discount rates: $\beta \in \{99.65\%, 99.60\%, 99.55\%\}$. The objective function is (5.3) in Section 5. We modify (I.1) to allow for risk aversion

$$u(y_t, s_{p,t}, s_{f,t}) = \pi U(y_t; \rho) - R(s_{p,t}, s_{f,t}; \rho),$$

where $U(y_t; \rho)$ represents either a risk-neutral agent ($\rho = 0$) or CRRA with risk aversion parameter ρ (e.g., $\frac{y^{1-\rho}-1}{1-\rho}$), and $R(\cdot, \cdot; \rho)$ is the regret function. The dataset is discretized over 400 points according to the procedure laid out in Section G.

The following utility function (Model 1) is estimated in Table I5

$$u(y_t, s_{p,t}, s_{f,t}; \rho) = \pi U(y_t; \rho) - \omega U(s_{p,t}; \rho) - \alpha U(s_{f,t}; \rho).$$

The first three columns refer to the linear utility case while the remaining part of the table reports estimates for the CRRA case.

Table I6 includes different coefficients for the two conditions and an interaction term as shown in (I.3).

⁹At this point we should mention that the results of the estimations are very similar in all models if we assume CRRA utility function instead of the linear one. In this case the regret term becomes $R(s_{p,t}, s_{f,t}; \rho) = \mathbb{1}_{\{\text{No Info}\}} \omega_{NI} U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} (\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t}; \rho) U(s_{f,t}; \rho))$, where ρ is a risk preference parameter in $U(y; \rho) = (y^{1-\rho} - 1)/(1 - \rho)$. The same estimation as in Table I4, only with an additional parameter ρ , is presented in Table I7 in Appendix I.2, which also contains several other model specifications. Overall, the estimated risk preferences are close to risk neutrality in all alternative models and the coefficients on the rest of the parameters stay similar.

Two models are reported with different interaction terms. In Model 2 the regret term is specified as

$$R = \mathbb{1}_{\{\text{No Info}\}} \omega_{NI} U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} (\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t} \times s_{f,t}; \rho))$$

while in Model 3 the regret term is:

$$R = \mathbb{1}_{\{\text{No Info}\}} \omega_{NI} U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} (\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I s_{p,t} \times s_{f,t}).$$

The last two equations assume CRRA utility because in the linear case they would produce the same estimate as those in the rightmost columns of Table I4.

Table I7 replicates the results of Table I4. The estimates show that the utility function is almost linear the regret parameters stay qualitatively unchanged.

Finally, Table I8 estimates $\pi U(y_t; \rho) - R(s_{p,t}, s_{f,t}; \rho)$ where the regret term is defined as

$$R(s_{p,t}, s_{f,t}; \rho) = \mathbb{1}_{\{\text{No Info}\}} (\omega_{NI} U(s_{p,t}; \rho) + \alpha_{NI} U(s_{f,t}; \rho) + \lambda_{NI} U(s_{p,t}; \rho) U(s_{f,t}; \rho)) \\ + \mathbb{1}_{\{\text{Info}\}} (\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t}; \rho) U(s_{f,t}; \rho))$$

in Model 4 and

$$R(s_{p,t}, s_{f,t}; \rho) = \mathbb{1}_{\{\text{No Info}\}} (\omega_{NI} U(s_{p,t}; \rho) + \alpha_{NI} U(s_{f,t}; \rho) + \lambda_{NI} U(s_{p,t} \times s_{f,t}; \rho)) \\ + \mathbb{1}_{\{\text{Info}\}} (\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t} \times s_{f,t}; \rho))$$

in Model 5, where $U(\cdot; \rho)$ is a CRRA utility function (the estimates report very mild risk-seeking preferences, while the regret parameters do not change substantially). Overall, the results are very similar across all tables, and corroborate our conclusions outlined in Section 7.

Section I.3 estimates a similar model allowing for loss aversion.

	Model 1					
	Linear Utility			CRRA Utility		
	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
$\hat{\rho}$				0.061*** (0.017)	0.061*** (0.017)	0.062*** (0.017)
$\hat{\pi}$	1.892*** (0.011)	1.890*** (0.011)	1.887*** (0.011)	2.068*** (0.052)	2.068*** (0.052)	2.069*** (0.052)
$\hat{\omega}$	0.128 (0.144)	0.213 (0.142)	0.313* (0.139)	0.235 (0.170)	0.308* (0.170)	0.400** (0.169)
$\hat{\alpha}$	0.176*** (0.076)	0.219** (0.075)	0.255*** (0.073)	0.122 (0.086)	0.161* (0.085)	0.196** (0.083)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table I5: The estimation of (5.3) with the regret terms as in Model 1, specified in Appendix I.2 in periods 16 to 48 for different values of the discount factor β . The CCP are defined as in (6.1), maintaining the coefficient for the future regret term (β_3) equal to zero in the No Info condition. Standard errors are in parenthesis.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

	Model 2			Model 3		
	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
CRRA Utility						
$\hat{\rho}$	0.070*** (0.017)	0.074*** (0.017)	0.072*** (0.017)	0.074*** (0.017)	0.075*** (0.017)	0.076*** (0.017)
$\hat{\pi}$	2.089*** (0.053)	2.094*** (0.053)	2.091*** (0.053)	2.099*** (0.054)	2.101*** (0.054)	2.103*** (0.054)
$\hat{\omega}_{NI}$	0.548** (0.225)	0.743*** (0.217)	0.662*** (0.221)	0.567** (0.227)	0.616*** (0.226)	0.683*** (0.223)
$\hat{\omega}_I$	2.234*** (0.588)	2.516*** (0.448)	2.441*** (0.484)	2.272*** (0.570)	2.368*** (0.511)	2.442*** (0.465)
$\hat{\alpha}_I$	1.965*** (0.450)	2.136*** (0.337)	2.098*** (0.364)	1.947*** (0.424)	2.008*** (0.376)	2.046*** (0.340)
$\hat{\lambda}_I$	-0.317*** (0.080)	-0.349*** (0.065)	-0.341*** (0.068)	-0.241*** (0.054)	-0.248*** (0.049)	-0.253** (0.045)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table I6: The estimation of (5.3) with the regret terms as in Models 2 and 3, specified in Appendix I.2 in periods 16 to 48 for different values of the discount factor β . The CCP are defined as in (6.1), maintaining the coefficient for the future regret term (β_3) equal to zero in the No Info condition. Standard errors are in parenthesis.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

	Estimation of (I.2)			Estimation of (I.3)		
	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
CRRA Utility						
$\hat{\rho}$	0.060*** (0.017)	0.060*** (0.017)	0.062*** (0.017)	0.070*** (0.017)	0.071*** (0.017)	0.073*** (0.017)
$\hat{\pi}$	2.066*** (0.052)	2.066*** (0.052)	2.068*** (0.052)	2.089*** (0.053)	2.090*** (0.053)	2.092*** (0.053)
$\hat{\omega}_{NI}$	0.454** (0.219)	0.499** (0.218)	0.567*** (0.215)	0.548** (0.225)	0.596** (0.224)	0.662** (0.221)
$\hat{\omega}_I$	-0.008 (0.224)	0.091 (0.223)	0.202 (0.221)	1.917*** (0.515)	2.020*** (0.464)	2.100*** (0.424)
$\hat{\alpha}_I$	0.120 (0.086)	0.169** (0.085)	0.212** (0.084)	1.647*** (0.374)	1.714*** (0.331)	1.757*** (0.300)
$\hat{\lambda}_I$				-0.295*** (0.072)	-0.307*** (0.066)	-0.317*** (0.061)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table I7: The estimation of (5.3) with the regret terms as in (I.2) and (I.3) in periods 16 to 48 for different values of the discount factor β . The CCP are defined as in (6.1), maintaining the coefficient for the future regret term (β_3) equal to zero in the No Info condition. Standard errors are in parenthesis.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

	Model 4			Model 5		
	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
	CRRA Utility					
$\hat{\rho}$	-0.134*** (0.022)	-0.133*** (0.022)	-0.131*** (0.022)	-0.134*** (0.022)	-0.133*** (0.022)	-0.131*** (0.022)
$\hat{\pi}$	1.474*** (0.048)	1.476*** (0.048)	1.478*** (0.048)	1.474*** (0.048)	1.476*** (0.048)	1.478*** (0.048)
$\hat{\omega}_{NI}$	0.431 (0.435)	0.633 (0.393)	0.784** (0.360)	0.415 (0.470)	0.635 (0.425)	0.799** (0.388)
$\hat{\omega}_I$	1.526*** (0.459)	1.601*** (0.415)	1.650*** (0.380)	1.670*** (0.499)	1.751*** (0.451)	1.803*** (0.413)
$\hat{\alpha}_{NI}$	-0.236 (0.314)	-0.075 (0.278)	0.042 (0.250)	-0.251 (0.351)	-0.073 (0.311)	0.057 (0.280)
$\hat{\alpha}_I$	1.188*** (0.329)	1.230*** (0.291)	1.252*** (0.262)	1.332*** (0.371)	1.380*** (0.329)	1.406*** (0.297)
$\hat{\lambda}_{NI}$	0.018 (0.043)	-0.002 (0.039)	-0.017 (0.036)	0.016 (0.038)	-0.002 (0.035)	-0.015 (0.032)
$\hat{\lambda}_I$	-0.163*** (0.048)	-0.170*** (0.044)	-0.174*** (0.040)	-0.144*** (0.043)	-0.150*** (0.040)	-0.153** (0.037)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table I8: The estimation of (5.3) with the regret terms as as in Model 4 and 5, specified in Appendix I.2 in periods 16 to 48 for different values of the discount factor β . The CCP are defined as in (6.1) for both conditions. Standard errors are in parenthesis.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

I.3 Loss Aversion

Table I9 reports estimates for the model in (6.2) allowing for loss aversion. Loss aversion is defined as the utility loss due to selling at a price below the entry price. The regret function including loss aversion, $\tilde{R}(s_{p,t}, s_{f,t})$ is defined as:

$$\mathbb{1}_{\{\text{No Info}\}}(\omega_{NI}s_{p,t} + \alpha_{NI}s_{f,t} + \lambda_{NI}s_{p,t} \times s_{f,t}) + \mathbb{1}_{\{\text{Info}\}}(\omega_I s_{p,t} + \alpha_I s_{f,t} + \lambda_I s_{p,t} \times s_{f,t}) + \psi \mathbb{1}_{\{y_t < \text{entry price}\}}(y_t - \text{entry price})$$

Table I10 shows the same analysis allowing for CRRA risk preferences. In both models the coefficient ψ multiplies the negative loss. Define $\tilde{R}(s_{p,t}, s_{f,t}; \rho, \psi)$ to include both regret and loss aversion as:

1. $\mathbb{1}_{\{\text{No Info}\}}(\omega_{NI}U(s_{p,t}; \rho) + \alpha_{NI}U(s_{f,t}; \rho) + \lambda_{NI}U(s_{p,t}; \rho) \times U(s_{p,t}; \rho))$
 $+ \mathbb{1}_{\{\text{Info}\}}(\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t}; \rho) \times U(s_{f,t}; \rho)) + \psi \mathbb{1}_{\{y_t < \text{entry price}\}}(y_t - \text{entry price})$
2. $\mathbb{1}_{\{\text{No Info}\}}(\omega_{NI}U(s_{p,t}; \rho) + \alpha_{NI}U(s_{f,t}; \rho) + \lambda_{NI}U(s_{p,t} \times s_{f,t}; \rho))$
 $+ \mathbb{1}_{\{\text{Info}\}}(\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t} \times s_{f,t}; \rho)) + \psi \mathbb{1}_{\{y_t < \text{entry price}\}}(y_t - \text{entry price})$

In these tables ψ is constrained to be non-negative, as participants would enjoy a loss otherwise.

Parameter	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
Linear Utility			
$\hat{\pi}$	1.832*** (0.022)	1.830*** (0.022)	1.829*** (0.022)
$\hat{\omega}_{NI}$	1.075*** (0.464)	1.255*** (0.416)	1.386*** (0.378)
$\hat{\omega}_I$	2.195*** (0.477)	2.240*** (0.427)	2.266*** (0.388)
$\hat{\alpha}_{NI}$	0.070 (0.341)	0.214 (0.303)	0.316 (0.274)
$\hat{\alpha}_I$	1.665*** (0.348)	1.680*** (0.309)	1.681*** (0.281)
$\hat{\lambda}_{NI}$	-0.030 (0.051)	-0.050 (0.046)	-0.064 (0.043)
$\hat{\lambda}_I$	-0.244*** (0.053)	-0.247*** (0.048)	-0.249*** (0.044)
$\hat{\psi}$	0.000 (0.024)	0.000 (0.024)	0.000 (0.024)
N	111,613	111,613	111,613

Table I9: Estimation of regret and loss-aversion parameter in the risk-neutral case. Periods: 16 to 48. Standard errors are in parenthesis. $\hat{\psi}$ is the estimated coefficient of loss aversion.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

Parameter	$\beta = 99.65\%$	$\beta = 99.60\%$	$\beta = 99.55\%$
Model 1			
$\hat{\rho}$	-0.051*** (0.024)	-0.053** (0.024)	-0.049** (0.024)
$\hat{\pi}$	1.698*** (0.070)	1.690*** (0.069)	1.699*** (0.070)
$\hat{\omega}_{NI}$	0.505 (0.492)	0.696 (0.441)	0.862** (0.405)
$\hat{\omega}_I$	1.664*** (0.517)	1.714*** (0.464)	1.768*** (0.426)
$\hat{\alpha}_{NI}$	-0.246 (0.355)	-0.079 (0.312)	0.033 (0.282)
$\hat{\alpha}_I$	1.323*** (0.371)	1.354*** (0.327)	1.364*** (0.296)
$\hat{\lambda}_{NI}$	0.019 (0.055)	-0.004 (0.050)	-0.020 (0.046)
$\hat{\lambda}_I$	-0.204*** (0.061)	-0.208*** (0.055)	-0.211*** (0.052)
$\hat{\psi}$	0.000 (0.027)	0.000 (0.027)	0.000 (0.027)
Model 2			
$\hat{\rho}$	-0.063** (0.024)	-0.051* (0.024)	-0.055** (0.024)
$\hat{\pi}$	1.667*** (0.069)	1.698*** (0.070)	1.683*** (0.069)
$\hat{\omega}_{NI}$	0.505 (0.529)	0.747 (0.486)	0.936** (0.440)
$\hat{\omega}_I$	1.520*** (0.557)	1.958*** (0.515)	2.107*** (0.469)
$\hat{\alpha}_{NI}$	-0.210 (0.398)	-0.045 (0.360)	0.102 (0.322)
$\hat{\alpha}_I$	1.254*** (0.417)	1.583*** (0.381)	1.679*** (0.343)
$\hat{\lambda}_{NI}$	0.011 (0.050)	-0.008 (0.048)	-0.026 (0.043)
$\hat{\lambda}_I$	-0.155*** (0.056)	-0.202*** (0.055)	-0.214*** (0.051)
$\hat{\psi}$	0.000 (0.027)	0.000 (0.027)	0.000 (0.027)
N	111,613	111,613	111,613

Table I10: Estimation of models 1 and 2. Periods: 16 to 48. Standard errors are in parenthesis. $\hat{\psi}$ is the estimated coefficient of loss aversion.

***, **, * denote statistical significance at the 1, 5 and 10 percent level.

J Instructions (English)

J.1 Market Task

J.1.1 General Instructions

Dear Participants,

You are participating in a decision making experiment which consists of a main part and a questionnaire. If you follow the instructions carefully, you can earn a considerable amount of money depending on your decisions and random events. Your earnings will be paid to you at the end of the experiment.

During the experiment you are not allowed to communicate with anybody. In case of questions, please raise your hand. Then we will come to your seat and answer your questions. Any violation of this rule excludes you immediately from the experiment and all payments.

In the end of the experiment the payment will be made in **CASH**.

J.1.2 The Task

In this experiment you will make decisions in 48 different tasks. Each task is separate and does not depend on the previous tasks in any way. At the beginning of each task you receive 10 Euro. You can earn or lose money depending on your choices. This money will be added or subtracted from 10 Euro.

J.1.3 Structure of the Task

Imagine that you are participating in a financial market and that you should decide at each market (trial) when to sell an object. At the beginning of each market (trial) you observe the price of an object for 15 periods (each period lasts 0.8 seconds). During these periods you can see how the price of the object evolves before you enter the market which means that you cannot make any decisions during these 15 periods. The picture on the right shows the example of the price of the object varying during this starting phase. When you see a vertical red line drawn across the graph, this means that the starting phase of price observation is over. The current price of the object at this point corresponds to the price at which you enter the market. On the top of the screen you can see the current price displayed in each period (between €0 and €10).



J.1.4 The Process Guiding the Value

In every market the value changes according to the following process. If the value in the current period is V , then the value in the next period depends on 1) the current value V and 2) the randomly generated number S . In particular, the value in the next period is equal to $0.6V + S$, where S is a number between 0 and 4. This means that in each period any number between 0 and 4 (for example, 2.1789 or 3.51) has equal probability of being chosen and will contribute to the future value. Therefore, any number in the interval between $0.6V$ and $0.6V + 4$ has equal probability to be the value of the object in the next period. The table below shows some examples. Notice also that in each period the current value cannot be higher than € 10 and lower than € 0.

CURRENT VALUE	INTERVAL FOR THE VALUE IN THE NEXT PERIOD	
	MINIMAL VALUE	MAXIMAL VALUE
€2	€1.2	€5.2
€4	€2.4	€6.4
€6	€3.6	€7.6
€8	€4.8	€8.8

J.1.5 Entering the Market

After you have observed the evolution of the value for 15 periods the market stops at the red vertical line and the button ENTRATA (ENTER) appears at the bottom of the screen (see top figure). When you press the button you enter the market. This means that you “buy” the object at the current value and spend €2.59 as indicated at the top of the screen. You do not have a choice at which price to buy the object. Once you press the button three things happen: 1) the Valore di entrata (Entry price) appears on top of the screen in red (see bottom figure); 2) the value starts to change again and 3) the button changes to USCITA (EXIT).



J.1.6 Exiting the Market

The choice you make in the market is when to exit. This is the point at which you “sell” the object and obtain the amount of money equal to the current value. Your profit in the market is the amount you received at the exit minus the amount you paid when you entered. For example, if you entered at the value of €2.59 and exited at the value of €2.68 your profit is $2.68 - 2.59 = 0.09$, or 9 cents. If you entered at the value of €2.59 and exited at the value of €2.45 your profit is $2.45 - 2.59 = -0.14$, or minus 14 cents. Thus, **YOUR PROFIT CAN BE NEGATIVE**. If you do not choose to exit before the closure of the market at period 50, your profit will be calculated using the last period value of the object.

J.1.7 Observed and Unobserved Future

There are two possible scenarios, which can happen after you press the USCITA (EXIT) button, or sell the object. In one scenario you will observe the evolution of the value of the object until the market closure (after period 50). In the other case you will not observe the evolution of the value. You will be informed about which scenario you are in **BEFORE** the opening of each market. Before each market you will observe a screen with two possible phrases: "INFO DOPO l'uscita" (Information after exit) or "NO INFO al'uscita" (No information after exit) (see figures). The former indicates that the market which you will choose in next is the one with observable future value and the latter – the market with non-observable future value. To make sure that you remember which scenario you are in, the "INFO DOPO" and "NO INFO" signs will appear in the top left corner of the screen while the market is evolving.

INFO DOPO l'uscita

NO INFO all'uscita

J.1.8 After Exiting the Market

After you exit the market, or press USCITA (EXIT) button, you will be provided with the information on your profit. Top figure illustrates the scenario with observable future and the bottom figure – with non-observable future. In both cases, you will see the “Valore di uscita” (exit value) in blue and profit in green (if positive) or red (if negative). In case of non-observable future you will be also asked to wait until the market closure which is the same time it would have taken the market to reach closure if you could have observed the future value. When the market closes you can press PROSEGUI (CONTINUE) button to proceed to the next market.



J.1.9 Payment

Your payment in the experiment is determined as follows. Before the experiment you are given an endowment of €10. After you finish choosing in all 50 markets, one of them will be chosen at random and the profit that you made in that market will be added to your endowment. So, if you earned €3 in the chosen market, your total payment will be €10 + €3 = €13. If your profit was -€3, your total payment will be €10 - €3 = €7. Notice that your profit can change between -€10 and €10. Thus you can earn minimum of €0 and maximum of €20.

J.1.10 Trial Markets

Before the beginning of the task you will have an opportunity to familiarize yourself with the interface in 6 trial markets which will look exactly the same as the actual markets but with TRIAL DI PROVA (TRIAL MARKET) written on the screen. **You will not be paid for your decisions in trial markets.**

K Instructions (Italian)

K.1 Market Task

K.1.1 Informazioni Generali

Gentile partecipante,

Prenderai parte ad un esperimento comprendente due compiti decisionali e un questionario. Se segui le istruzioni attentamente potrai guadagnare una considerevole somma di denaro, che dipenderà dalle decisioni che prenderai durante l'esperimento. La somma da te guadagnata ti verrà pagata al termine dell'esperimento.

Ti chiediamo per favore di non comunicare con gli altri partecipanti durante l'esperimento. Nel caso tu abbia delle domande, chiedi allo sperimentatore alzando la mano. A quel punto lo sperimentatore verrà alla tua postazione e risponderà alle tue domande.

Al termine dell'esperimento il pagamento verrà effettuato in **CONTANTI**.

K.1.2 Compito di Scelta

In questo compito ti verrà chiesto di prendere una decisione in 48 diversi problemi. Ogni problema è a se stante e non dipende dall'esito ottenuto nei problemi precedenti. All'inizio del compito riceverai una somma di partenza pari a 10 euro. In ogni problema potrai guadagnare o perdere un certo ammontare di denaro, il quale verrà sommato o sottratto a questi 10 euro.

K.1.3 Struttura di Campito di Scelta

Immagina di essere all'interno di un mercato finanziario e di dover decidere, ad ogni trial, quando incassare l'ammontare investito. Ogni mercato (trial) inizia osservando il valore dell'oggetto del tuo investimento per 15 periodi (ogni period dura 0.8 secondi). Durante questa prima fase, vedrai come il valore dell'oggetto si è evoluto nei precedenti 15 periodi del mercato. Durante questi 15 periodi non potrai prendere nessuna decisione. La figura a destra ti mostra un esempio di come il valore dell'oggetto può variare durante questa prima fase. Quando la linea verticale rossa verrà raggiunta, significa che i 15 periodi della fase di osservazione saranno terminati. A quel punto il valore corrente dell'oggetto corrisponderà al tuo valore d'entrata nel mercato. La dicitura **"Valore corrente"** in alto ti mostra il valore dell'oggetto in ogni periodo (tra €0 e €10).



K.1.4 Il Processo Che Stabilisce il Valore

In ogni mercato il prezzo cambia seguendo un particolare processo. Dato il valore corrente in un periodo del mercato, V , il valore nel periodo successivo (all'interno dello stesso mercato) dipende da 1) il valore corrente, V , e 2) un numero generato in maniera random, S . In particolare, il valore nel periodo seguente è uguale a $0.6V + S$, dove S è un numero tra 0 e 4. Ciò significa che in ogni periodo qualunque numero tra 0 e 4 (per es. 2.1789 o 3.51) ha la stessa probabilità di essere scelto e di contribuire al valore futuro. Perciò ogni numero nell'intervallo tra $0.6V$ e $0.6V + 4$ ha la stessa probabilità di essere il valore dell'oggetto nel prossimo periodo. La tabella qui di seguito riporta alcuni esempi. Nota che in ogni periodo il valore corrente non può essere maggiore di €10 né minore di €0.

VALORE CORRENTE	INTERVALLO DEL VALORE NEL PERIODO SUCCESSIVO VALORE MINIMO	VALORE MASSIMO
€2	€1.2	€5.2
€4	€2.4	€6.4
€6	€3.6	€7.6
€8	€4.8	€8.8

K.1.5 Entrare nel Mercato

Dopo aver osservato 15 periodi il mercato si fermerà alla linea verticale rossa e il pulsante “ENTRATA” apparirà in basso (vedi la figura in alto a destra). A questo punto per entrare nel mercato dovrai premere il tasto “ENTRATA.” Questo significa che effettivamente tu compri l’oggetto al valore corrente. Nell’esempio indicato nella figura in alto spenderesti €2.59. Non ti sarà possibile evitare di entrare nel mercato e non potrai scegliere tu stesso a quale prezzo comprare l’oggetto. Una volta premuto il pulsante “ENTRATA” il valore dell’oggetto comincerà a variare nuovamente e ti compariranno tre nuove informazioni a schermo (figura in basso a destra): 1) il “Valore di entrata” in rosso in alto a sinistra; 2) il valore attuale dell’oggetto; 3) il pulsante “USCITA.”



K.1.6 Uscire dal Mercato (Uscita)

L’unica scelta a tua disposizione in ogni mercato sarà quando uscire. Questa scelta corrisponde al momento in cui decidi di vendere l’oggetto e intascare la somma di denaro pari al “Valore corrente.” Il tuo guadagno nel mercato sarà la differenza tra il “Valore corrente” al momento di vendita dell’oggetto e il “Valore di entrata.” Ad esempio, se tu entri quando l’oggetto vale €2.59 ed esci al valore di €2.68 il tuo guadagno sarà pari a $€2.68 - €2.59 = €0.09$, o 9 centesimi. Se invece entri al “Valore di entrata” pari a €2.59 ed esci quando il “Valore corrente” è €2.45, il tuo guadagno sarà di $€2.45 - €2.59 = €-0.14$, o un guadagno negativo di 14 centesimi. Perciò, **IL TUO GUADAGNO NEL MERCATO PUO’ ESSERE NEGATIVO.** Se non esci prima della fine del mercato, che dura 50 periodi, il tuo guadagno sarà calcolato usando il valore corrente nell’ultimo periodo.

K.1.7 Futuro Osservato o non Osservato

Ci sono due possibili scenari alternativi che si possono realizzare dopo che hai cliccato sul pulsante “USCITA,” ovvero venduto l’oggetto. In uno scenario ti verrà mostrata l’evoluzione del valore dell’oggetto fino alla chiusura del mercato (50esimo periodo). Nell’altro caso, dopo la vendita dell’oggetto non osserverai nulla, e un nuovo mercato inizierà. Sarai informato riguardo allo scenario in cui ti trovi **PRIMA** dell’inizio di ogni mercato. Prima di ogni mercato, osserverai una schermata con due possibili frasi: “**INFO DOPO l’uscita**” o “**NO INFO all’uscita**” (vedi le figure a destra). La prima dicitura indica che ti trovi in un mercato in cui l’evoluzione del valore dopo la vendita è osservabile, mentre la seconda dicitura ti informa che il futuro valore dell’oggetto non è osservabile. Per ricordarti in quale scenario ti trovi, le diciture “INFO DOPO” e “NO INFO” sono mostrate in alto a sinistra della schermata in cui vedi l’evoluzione del mercato.

INFO DOPO l'uscita

NO INFO all'uscita

K.1.8 Dopo Essere Usciti dal Mercato

Dopo la tua uscita dal mercato, o dopo aver premuto il pulsante “USCITA,” riceverai informazioni sul tuo guadagno. La figura in alto a destra ti mostra lo scenario “INFO DOPO,” dove il futuro è osservabile, mentre la figura in basso ti mostra lo scenario “NO INFO,” dove il futuro non è osservabile. In entrambi i casi, in alto a destra visualizzerai il “Valore di uscita” in blu, ed il tuo “Guadagno” in verde se positivo e in rosso se negativo. Inoltre, nello scenario Info Dopo dovrai attendere il termine del mercato, che corrisponde al tempo che il mercato avrebbe impiegato per raggiungere la sua naturale conclusione (50 periodi) se tu non avessi venduto l’oggetto prima. Raggiunto l’ultimo periodo potrai esaminare la tua prova; per accedere al prossimo mercato dovrai cliccare sul pulsante “Prosegui.”



K.1.9 Pagamento

Il tuo guadagno nell’esperimento viene calcolato come segue. Prima dell’esperimento ti vengono dati € 10 a disposizione. Quando hai finito di scegliere in tutti i 48 mercati, uno di questi verrà scelta in modo casuale e il guadagno che tu fai in quel mercato sarà sommato ai € 10 di partenza. Perciò, se tu guadagni € 3 nel mercato scelto, il tuo pagamento totale sarà € 10 + € 3 = € 13. Nel caso di un guadagno negativo, ad esempio -€ 3, il tuo pagamento totale sarà € 10 - € 3 = € 7. Nota che il tuo guadagno può variare tra -€ 10 e +€ 10, perciò il tuo pagamento totale varia tra un minimo di € 0 e un massimo di € 20.

K.1.10 Mercati di Prova

Prima dell’inizio del compito ti viene data l’opportunità di familiarizzare con l’interfaccia in 6 mercati di prova che assomigliano in tutto e per tutto ai mercati reali a cui parteciperai successivamente, con l’unica differenza che in questi mercati la dicitura TRIAL DI PROVA compare sullo schermo. **Non verrai pagato per le tue decisioni nei mercati di prova.**

K.2 Holt and Laury Task (Italian)

DESCRIZIONE DELLA SECONDA PARTE DELL'ESPERIMENTO

In questa parte dell'esperimento ti verranno presentate 10 coppie di lotterie. Ogni lotteria ti garantisce di ottenere, con una certa probabilità, una tra due possibili vincite. Per ogni coppia di lotterie, il tuo compito sarà quello di scegliere la lotteria che preferisci giocare. Di seguito ti verrà presentata una descrizione dettagliata del compito. Premere il pulsante OK per continuare.

DESCRIZIONE DEL COMPITO

Nella parte destra dello schermo sono riportate le 10 coppie di lotterie. Ci sono 10 righe che corrispondono alle 10 scelte che dovrai effettuare. Ogni riga rappresenta una scelta tra due lotterie.

Per effettuare le tue scelte sarà sufficiente cliccare in corrispondenza della lotteria che preferisci. Una volta che avrai scelto una lotteria, essa diventerà di colore rosso.

Dopo che avrai effettuato le tue 10 scelte, il computer selezionerà in modo casuale una delle 10 righe. Infine, la lotteria da te scelta verrà giocata dal computer e tu riceverai la vincita corrispondente all'esito della lotteria. La tua vincita ti verrà mostrata a schermo dopo che avrai completato e validato tutte le tue scelte.

Ricorda, l'ammontare di denaro rappresentato nelle diverse lotterie è reale, perciò sarai pagato/a in base alle scelte che effettuerai e secondo le regole appena descritte.

Se hai qualche dubbio sulla procedura ed il metodo di pagamento sentiti libero/a di chiedere chiarimenti allo sperimentatore.

1	10% prob. vincere €2.30 90% prob. vincere €1.60	OPPURE	10% prob. vincere €4.00 90% prob. vincere €0.20
2	20% prob. vincere €2.30 80% prob. vincere €1.60	OPPURE	20% prob. vincere €4.00 80% prob. vincere €0.20
3	30% prob. vincere €2.30 70% prob. vincere €1.60	OPPURE	30% prob. vincere €4.00 70% prob. vincere €0.20
4	40% prob. vincere €2.30 60% prob. vincere €1.60	OPPURE	40% prob. vincere €4.00 60% prob. vincere €0.20
5	50% prob. vincere €2.30 50% prob. vincere €1.60	OPPURE	50% prob. vincere €4.00 50% prob. vincere €0.20
6	60% prob. vincere €2.30 40% prob. vincere €1.60	OPPURE	60% prob. vincere €4.00 40% prob. vincere €0.20
7	70% prob. vincere €2.30 30% prob. vincere €1.60	OPPURE	70% prob. vincere €4.00 30% prob. vincere €0.20
8	80% prob. vincere €2.30 20% prob. vincere €1.60	OPPURE	80% prob. vincere €4.00 20% prob. vincere €0.20
9	90% prob. vincere €2.30 10% prob. vincere €1.60	OPPURE	90% prob. vincere €4.00 10% prob. vincere €0.20
10	100% prob. vincere €2.30 0% prob. vincere €1.60	OPPURE	100% prob. vincere €4.00 0% prob. vincere €0.20

OK

References

- ABBRING, J. H. (2010). Identification of dynamic discrete choice models. *Annual Review of Economics*, **2** (1), 367–394.
- AGUIRREGABIRIA, V. and MAGESAN, A. (2016). Solution and estimation of dynamic discrete choice structural models using Euler equations, Mimeo, University of Toronto and University of Calgary.
- and MIRA, P. (2010). Dynamic discrete choice structural models: a survey. *Journal of Econometrics*, **156** (1), 38–67.
- EIJKELENBOOM, G., ROHDE, I. and VOSTROKNUTOV, A. (2016). The impact of the level of responsibility on choices under risk: The role of blame, Mimeo, Maastricht University and University of Trento.
- FISCHBACHER, U. (2007). z-tree: Zurich Toolbox for ready-made economic experiments. *Experimental Economics*, **10** (2), 171–178.
- HOLT, C. A. and LAURY, S. K. (2002). Risk aversion and incentive effects. *The American Economic Review*, **92** (5), 1644–1655.
- MAGNAC, T. and THESMAR, D. (2002). Identifying dynamic discrete decision processes. *Econometrica: Journal of the Econometric Society*, **70** (2), 801–816.
- POTUSCHAK, H. and MULLER, W. G. (2009). More on the distribution of the sum of uniform random variables. *Statistical Papers*, **50** (1), 177–183.
- TAUCHEN, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, **20**, 177–181.