

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

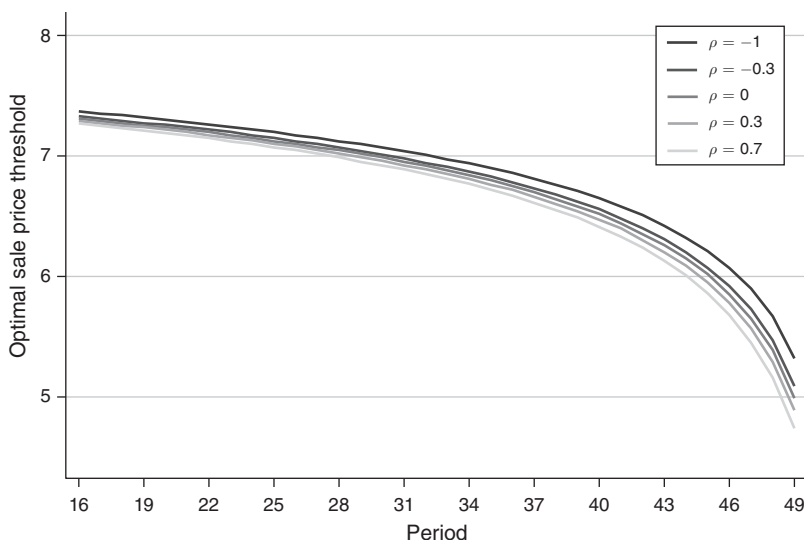


FIGURE 2. OPTIMAL SELLING PRICE THRESHOLDS

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.

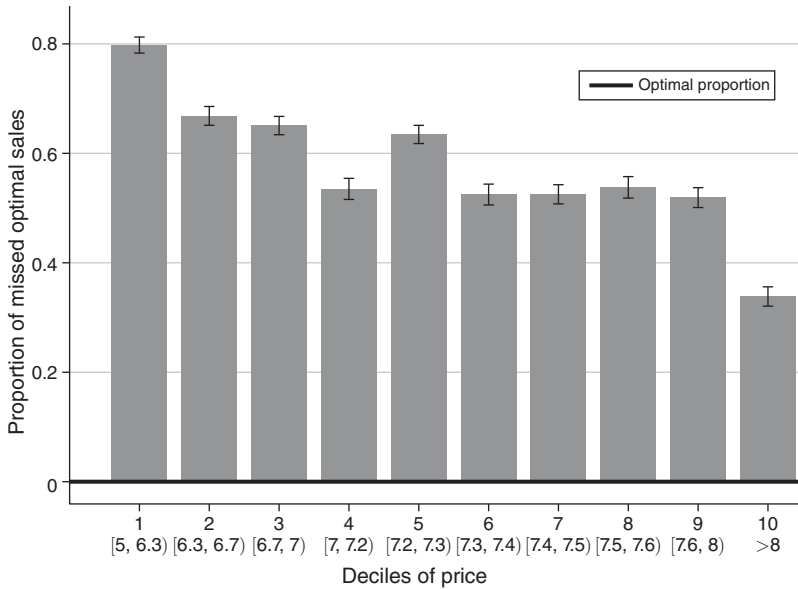


FIGURE 3. THE PROPORTION OF TIMES THE PARTICIPANTS DECIDED TO KEEP THE ASSET

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.

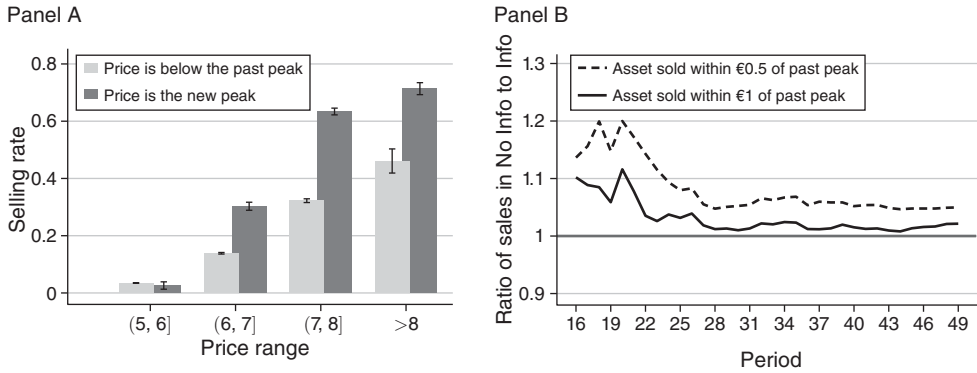


FIGURE 4. EVIDENCE OF REGRET AVOIDANCE

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).

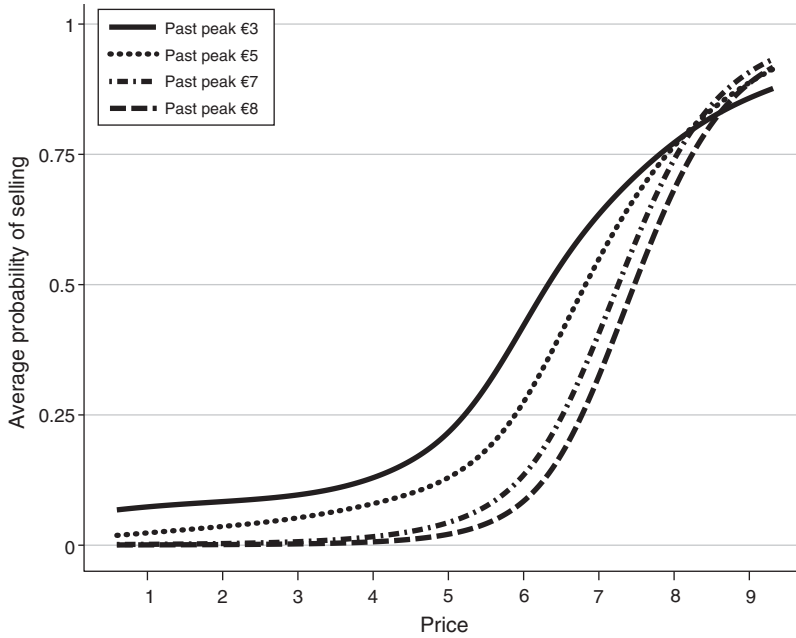


FIGURE 5. PROBABILITY OF SELLING THE ASSET IN THE NO INFO CONDITION

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

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

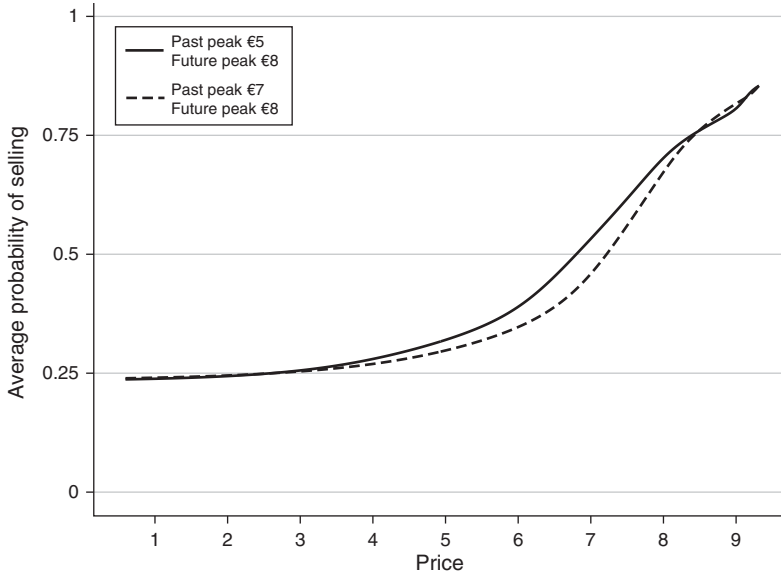


FIGURE 6. PROBABILITY OF SELLING THE ASSET IN THE INFO CONDITION

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.

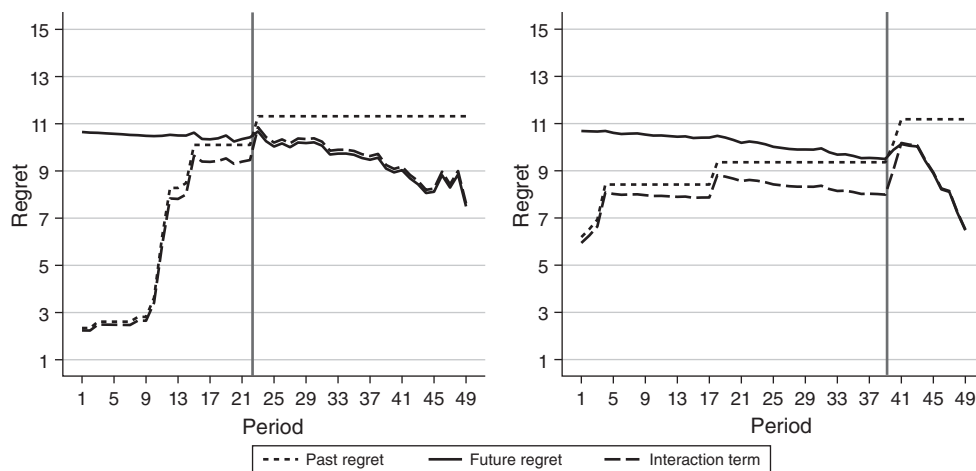


FIGURE 7. DYNAMICS OF PAST AND FUTURE REGRET

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=oqx49q>.
- 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.