Dynamic Regret Avoidance*

Michele Fioretti† ‡ **Alexander Vostroknutov**§ **Giorgio Coricelli**¶

July 2017

Abstract

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. The estimation of a structural model shows that without future prices participants are affected only by past regret, but when future prices are available, they also avoid future regret. Moreover, as the relative sizes of past and future regret change with the price, participants dynamically switch between them, with the larger regret term dominating the decision to sell.

JEL classifications: C51, C91, D03
Keywords: regret avoidance, dynamic regret, dynamic discrete choice, structural models, stock market behavior, experiments

*We are very much indebted to Geert Ridder for his insights and support in this project. We would also like to thank Martina Puppi and Folco Panizza for invaluable help with running the experiments. 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, Tasha Poppa, Alejandro Robinson Cortes, Joel Sobel, Jorge Tamayo, Michael Woodford, Leeat Yariv, Anthony Lee Zhang and other participants of seminars at CIMeC, Caltech, GMU and USC and at the 2017 AMES, 2017 ESA World Meeting, 2017 NAMS, M-BEES 2017 conferences. 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 ours.

†Department of Economics, University of Southern California, 3620 South Vermont Ave, Los Angeles, CA 90089. e-mail: fioretti@usc.edu

‡Corresponding author.

§Center for Mind/Brain Sciences, University of Trento, Via delle Regole 101, 38123 Mattarello (TN), Italy. e-mail: a.vostroknutov@unitn.it

¶Department of Economics, University of Southern California, 3620 South Vermont Ave, Los Angeles, CA 90089 and Center for Mind/Brain Sciences, University of Trento, Via delle Regole 101, 38123 Mattarello (TN), Italy. e-mail: giorgio.coricelli@usc.edu

1 Introduction

Regret is a negative emotion associated with an action or inaction which is felt 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 *et al.*, 2017; Frydman and Camerer, 2016), portfolio and pension scheme selection (Muermann *et al.*, 2006; Hazan and Kale, 2015), and currency hedging (Michenaud and Solnik, 2008).

Apart from the empirical applications, regret avoidance has been studied both theoretically (Savage, 1954; Bell, 1982; Loomes and Sugden, 1982; Skiadas, 1997; Sarver, 2008; Hayashi, 2008; Bikhchandani and Segal, 2014; Leung and Halpern, 2015; Qin, 2015) and experimentally (Coricelli *et al.*, 2005; Camille *et al.*, 2004; Zeelenberg and Beattie, 1997; Zeelenberg *et al.*, 1996; Bleichrodt *et al.*, 2010; Strack and Viefers, 2015). This research has mostly concentrated on *static* problems where a single decision is made that can be affected by the information about possible counterfactual outcomes. Many important real life decisions, like buying a house or a pension plan, fit into this setting.

Nevertheless, many interesting phenomena that involve regret have a *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 is taken today that brings about an outcome that was foregone in the past. Future regret involves taking actions that prevent missing the opportunity of achieving a more desirable expected future outcome. For example, in the financial markets the decision to sell an asset might depend on the highest observed price in the past, but traders might also think about the hypothetical counterfactual situation in which they sell today and regret it later, and adjust their behavior accordingly.

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 addressing 1) How strongly does the avoidance of past and future regret influence the choice to sell? 2) Is there an interaction between past and future regret? Does one become stronger or weaker in the presence of the other? 3) How do risk preferences interact with regret avoidance?

In our experiment, reminiscent of those in Oprea *et al.* (2009), Oprea (2014), and Strack and Viefers (2015), participants choose 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 if access to the prices after selling has an effect on the decision making (future regret). More importantly, our design makes it possible to use structural modelling 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 our 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 sell the asset to make the effect of past regret smaller or absent confirms the results of the recent studies which focus on past regret only (Gneezy, 2005; Strack and Viefers, 2015). We go further and show that agents should keep the asset longer when they expect high future maximum price, as compared to the no future regret case. 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, individual risk preferences do 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.

The effects of the past and future regret are not simply additive. With the use of a dynamic discrete choice model, we recover the participants' preferences and show the presence of an interaction between past and future regret in the utility function, which would not be possible to identify with simple regression analysis. We find that 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 which is stronger. Our findings demonstrate that individuals incorporate past and future regret into the utility function in dynamic settings, thus, showing that they are able to extract and update complex counterfactual information about the changing environment and integrate it into the decision process.

2 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 seconds intervals and had to decide when to sell an "asset" (see Figure 1). For the first 15 periods participants could only observe the price. 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 the 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).

In each market the price followed a stochastic mean reverting process defined by $y_{t+1} = \alpha y_t + (1-\alpha)\varepsilon$, where $\alpha = 0.6$, y_t is the price in period t, and ε 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.

Figure 1: Screenshots of two markets. 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 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."

Each participant made selling decisions in 48 different markets, which could be of two types. 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 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 onwards (INFO DOPO means "info after").

The markets were presented in random order that was generated independently for each participant. Half of the markets were presented in the No Info and half in the Info condition. The sequence of conditions was also randomized. After the markets, the participants were presented with incentivized Holt-Laury task (Holt and Laury, 2002) and a questionnaire. Overall, 154 participants took part in the experiment in 9 sessions. The experiment was programmed in z-Tree (Fischbacher, 2007). Further details of the design can be found in Appendix A.

3 Past and Future Regret

The evolution of the price allows us to estimate past and future regret in each period before participants make a selling decision. As in Gneezy (2005), Baucells *et al.* (2011), and Strack and Viefers (2015), we hypothesize that the highest price in the past is a *reference point* that our participants use to measure how well they are doing. This is a dynamic variable that changes when the price gets above the observe highest peak. Notice that without regret the optimal policy is to sell the asset whenever the price rises above €5 (for risk neutral agent) since the process governing the price is mean-reverting. Thus, in the no regret case the selling decision is independent of any reference points and past price history. In the regret case we conjecture that the higher the past peak the more negative the feeling of past regret should become given the current price (which is always less than or equal to the highest past peak). This implies that the past peak should exert influence on the decision to keep the asset. Similarly, if participants are aware that they will observe prices even after selling the asset, they can anticipate a situation when the future price will exceed the selling price, which would lead to negative emotions (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 market closure.[1] When this information is not available, future regret should not play any role in the selling decisions since it is known that the future price will not be observed. Past and future regret allow us to extend the notion of regret that is commonly used in static settings to dynamic environments. This allows us 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 model regret as an increasing function of the observed highest price in the past and the expected future highest price. This decision is motivated by recent work (e.g., Gneezy, 2005; Strack and Viefers, 2015) that leverages on the saliency of the past peak as the key measure of past regret, allowing us to disregard other quantiles of the distribution of the past peak. Similarly,

[1] Sensitivity to future regret is also documented in Strahilevitz *et al.* (2011) who found that investors are less likely to repurchase an asset, which they previously sold and which gained value, as they are "painfully aware" of their loss of money by selling too early. This phenomenon is explained by the idea that participants try to detach themselves from a previously sold winning asset to avoid being automatically reminded of their poor judgement. Even though our experiment excludes repurchase, this event can be classified as a realization of future regret avoidance.

we model future regret as the expectation of the highest possible price achievable in any of the remaining future periods, given the current price and the number of periods left.[2]

In this section we define past and future regret and provide intuition on how they should affect the optimal selling decision. Note that the definitions given in this section are only to illustrate the main idea. The estimated model is described in Section 4. Under past regret (only), the agent incurs a disutility from not having sold at the past maximum. Suppose her utility in this case is $U(y_t) - U(s_{p,t})$, where $s_{p,t} = \max_{s \le t}\{y_s\}$. Since, by definition, $U(y_t) - U(s_{p,t}) \le 0$, the highest attainable utility is 0, which implies that the agent should sell the asset if she reaches the past maximum. Formally, an agent sells if the utility at time t is larger than the maximum expected utility available in the remaining periods:

$$U(y_t) - U(s_{p,t}) \ge \max\{\mathbb{E}_{y_{t+1}}[U(y_{t+1}) - U(s_{p,t+1})|y_t], \mathbb{E}_{y_{t+1}}[v_{t+2}|y_t]\} \tag{3.1}$$

where $v_{t+2} = \max\{\mathbb{E}_{y_{t+2}}[U(y_{t+2}) - U(s_{p,t+2})|y_{t+1}], \mathbb{E}_{y_{t+2}}[v_{t+3}|y_{t+1}]\}$ and the value function in the last period T is $v_T = \mathbb{E}_{y_T}[U(y_T) - U(s_{p,T})|y_{T-1}]$.[3] In the experiment the price evolution is described by a Markov chain, and, thus, all expectations are conditional on the past price. Notice that the max operator is less than or equal to 0. Thus, the optimal policy is to sell in some vicinity of the current past maximum. Moreover, as $U(y_t)$ approaches $U(s_{p,t})$ from below, the difference between the utilities in the LHS and RHS of inequality (3.1) becomes larger. Thus, assuming random utility à la McFadden (1974), we should expect that the probability of selling increases in $U(y_t) - U(s_{p,t})$. This gives us the first prediction:

Prediction 1. *The probability of selling the asset is higher the closer the price is to the value of the past peak.*

It should be mentioned that $U(y_t) - U(s_{p,t})$ is a very specific utility function. When fitting the structural model we assume more generally that the utility is $U(y_t) - \kappa U(s_{p,t})$, as in Strack and Viefers (2015), where κ is a regret sensitivity parameter. In this case, qualitatively, Prediction 1 should stay the same as long as κ is in the vicinity of 1. When $\kappa = 0$ the optimal policy for the agent is to sell the asset whenever its price is above some threshold (€5 in case $U(y)$ is linear) which is independent of the past peak (see Appendix B). This means that the less the agent cares about past regret, the less his selling price is influenced by the past peak.[4]

The optimal policy becomes more complex when the agent can also experience future regret, which is the disutility from observing a price higher than the selling price at a future time point,

[2] Baucells *et al.* (2011) suggest that agents substitute old reference points with new ones as new information comes in, thus, we exclude prices other than the highest from the regret function.

[3] By design the participants in the last period are forced to sell at the current price. That is, no decision is taken at $t = T$.

[4] It should be acknowledged that this is not a standard regret aversion function which 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.

given that this information is available. Notice that the agent cannot be affected by future regret in isolation from past regret because the past peak is always revealed. Thus, we can detect future regret avoidance by comparing two conditions, No Info, where information about future price is not available and Info, where it is.

The expectation of the highest future peak at time t, denoted by $s_{f,t}$, is a function of the price today and the number of periods left until the market closure. $s_{f,t}$ is the expectation of the maximum price achievable in $T - t$ periods given the current price y_t. This quantity is increasing in y_t and decreasing in t because of the presence of a terminating period (see Appendix C for calculations). Conversely, notice that $s_{p,t}$ is a weakly increasing function of time, since it is defined as a maximum of the past prices. This suggests that in different points in time either the past or the future regret term becomes more dominant.

In the future regret case the agent sells if

$$U(y_t) - U(s_{p,t}) - U(s_{f,t}) \geq \max\{\mathbb{E}_{y_{t+1}}[U(y_{t+1}) - U(s_{p,t+1}) - U(s_{f,t+1})|y_t], \mathbb{E}_{y_{t+1}}[v_{t+2}|y_t]\} \tag{3.2}$$

where v_{t+2} now also includes the expectation over the future peak. With the idea of random utility in mind, the following predictions can be tested.

Prediction 2. *The probability of selling the asset at any fixed price level increases with time.*

Prediction 3. *The probability of selling the asset at any fixed price level decreases with the expected future peak.*

Prediction 2 follows from the fact that $s_{f,t}$ decreases in time and Prediction 3 from the expected future peak increasing in the current price.

Note that the predictions implied by (3.1) and (3.2) with respect to past and future regret hold for any weakly increasing utility function. In estimating the structural model below we focus on the class of CRRA utilities $U(x) = \frac{x^{1-\rho}-1}{1-\rho}$. Clearly, the optimal selling decision depends not only on past and future regret, but also on the risk aversion parameter. In Appendix B we show 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:

Prediction 4. *Other things equal, the probability of selling the asset increases in the degree of risk aversion.*

The predictions that we make in this section just describe the general directions in which past and future regret can influence the choice to sell. However, they do not allow us to *quantify* the parameters of the regret utility function that would truly let us understand how regret enters into the selling decision. In the next section we show how a structural model allows us to do that.

4 A Structural Model of Dynamic Regret Avoidance

The decision problem involves comparing the utility that the agent can obtain by selling in the current period with that expected by a future sale. Notice that this expectation does not simply depend on the current state described by the price, the past peak, and the expected future peak, but incorporates the Markov nature of the price dynamics and, more importantly, the optimal decisions that the agent might make in all possible futures.

The nature of the problem, therefore, dictates the use of a structural approach that would allow us to recover the regret averse utility function employed by the participants in our experiment.[5] We estimate a dynamic discrete choice model (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 6. Appendix E provides the full derivation of the model.

Analogously to the static discrete choice literature, in a dynamic environment participants choose their best action according to a threshold rule which, however, also takes into account the Markovian nature of the prices. In each period t a choice is made whether to sell the asset ($d = 0$) or to keep it ($d = 1$). $u^d(x_t)$ denotes the per period regret averse utility from choosing action d when the current state is x_t (the utility function will be parameterized in Section 6). A participant's intertemporal utility is

$$\mathbb{E}\left\{ \sum_{t=1}^{T} \beta^{t-1} u^d(x_t) \right\}$$

where the expectation is taken with respect to the future values of the independent variable $x_t = (y_t, s_{p,t}, s_{f,t})$ with y_t being the price in period t, $s_{p,t}$ the past peak, and $s_{f,t}$ the expected future peak. $\beta \in (0, 1)$ is the discount factor. The dynamic environment can be summarized using a value function v^d, representing the discounted sum of future payoffs and defined by a Bellman equation

$$v^d(x_t) = \begin{cases} 0 + \beta\mathbb{E}\{v(x_{t+1})|x_t, d = 1\} & \text{if } d = 1 \text{ (keep)} \\ u^0(x_t) + 0 & \text{if } d = 0 \text{ (sell).} \end{cases} \tag{4.1}$$

[5] Structural models to recover preferences from experimental data were used in the past. Bajari and Hortaçsu (2005) structurally estimated different behavioral models to rationalize bidding in experimental auctions and found that risk aversion provides the best explanation over a set of competing models. More recently, structural methods have been developed to identify and estimate level-k reasoning (e.g, Crawford and Iriberri, 2007; Gillen, 2009) and regret aversion (Hayashi and Yoshimoto, 2016) in auctions in the lab.

Notice that the per period payoff of continuing $u^1(x_t)$ and the continuation value of selling the asset are zero. In fact, participants are only paid the price at which they sell the asset at the time when they sell it. Similarly to most of the binary static discrete choice models, it is assumed that the value of each choice includes an additive logit error $(\varepsilon_t^1, \varepsilon_t^0)$. $v(x_{t+1}) = \int_\varepsilon \max\{v^0(x_{t+1}) + \varepsilon_{t+1}^0, v^1(x_{t+1}) + \varepsilon_{t+1}^1)\}\mathrm{d}\Lambda(\varepsilon)$ is the expected value function in the next period and takes this error into account.

To solve the Bellman equation 4.1 we show that the distribution of the observed choices uniquely identifies the utility function. Intuitively, the participant 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) showed that the per period utility is a function of these probabilities, which results in 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 $p^1(x_t) = \Pr(d = 1|x_t)$ for the probability of continuing and $p^0(x_t) = \Pr(d = 0|x_t)$ for the probability of selling. Importantly, the CCP can be estimated directly from the data. The identification procedure uses the CCP, together with the properties of the logit distribution, to simplify the Bellman equation in terms of known variables.

The logit assumption gives an analytical solution for the probability of choosing each action. For example, the probability of choosing action 1 is $p^1(x_t) = 1/\big(1 + \exp(v^0(x_t) - v^1(x_t)\big)$ which depends on the difference of the two alternative specific value functions (equation 4.1). This difference is

$$v^1(x_t) - v^0(x_t) = -u^0(x_t) + \beta \int_{\mathcal{X}_{t+1}} \int_\varepsilon \max\{v^0(x_{t+1}) + \varepsilon_{t+1}^0, v^1(x_{t+1}) + \varepsilon_{t+1}^1)\}\mathrm{d}\Lambda(\varepsilon)\mathrm{d}F(x_{t+1}|x_t) \tag{4.2}$$

where the inner integral is over the error term, and the outer one is over the state space in the next period. The law of motion of x_t across periods is characterized by the transition matrix $F(x_{t+1}|x_t)$ which describes the probability of moving from x_t to x_{t+1}. Note that since the only random variable is the current price, the transition matrix depends only on y_t.

The last equation can be simplified further. As noted by Hotz and Miller (1993) the CCP can be inverted resulting in $v^1(x_t) - v^0(x_t) = \ln(p^1(x_t)) - \ln(1 - p^1(x_t))$. This means that the left hand side of (4.2) is a known function of the data (the CCP). Let us denote it by $\phi(p^1(x_t))$ and rewrite equation (4.2) as

$$\phi(p^1(x_t)) = -u^0(x_t) + \beta \sum_{\mathcal{X}_{t+1}} \big(u^0(x_{t+1}) - \ln(p^0(x_{t+1}))\big) f(x_{t+1}|x_t). \tag{4.3}$$

where the summation substitutes the integration as the state space needs to be discrete (a technical requirement).[6] Note also that in passing from (4.2) to (4.3) the expectation of the optimal choice in the period $t+1$ is rewritten in terms of the utility from the terminating action (sale) and the probability of selling the asset in the next period. Appendix E shows all the steps of this derivation which is based on the properties of the logit error. Therefore, 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 (4.2). This gives us the objective function (equation 4.3) used to estimate a parameterized version of the utility of selling the asset $u^0(x_t)$, which includes both risk and regret averse components for the two conditions (Info and No Info). 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 on the Markovian structure of the state variables.[7]

5 Descriptive Statistics

In this section we test the predictions outlined in Section 3 with data summary statistics and regression analysis. These methods can provide only crude estimates of how the current state influences the choices to sell. Nevertheless, we find them instructive in showing how our theoretical intuition works. We start with comparisons of the behavior of our participants with the optimal choice of a regret free agent who should sell the asset whenever the price is above €5 if he is risk neutral (in case of risk aversion/seeking the threshold price decreases/increases, see Appendix B). We analyse the number of times that our participants could have sold the asset at a price above €5. Participants did not sell the asset at the price above €5 on average 6.99 times (SE: 0.11) in each market. This already shows that the behavior is not consistent with the regret free risk neutral agent who would have sold immediately after the price rose above €5. If we regress the average number of times of staying above €5 on the measure of risk aversion from the Holt-Laury task, we find that the most risk loving participants stay at the price above €5 for 1.6 periods longer than the most risk averse participants (coefficient on hl is -1.62, $p = 0.024$). This is consistent with Prediction 4 above: risk averse participants should sell at lower prices

[6] The 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 points is 0.049. This method is described in detail in Appendix D. The empirical results are robust to estimation on a finer grid although the optimization becomes more demanding (Appendix J.2).

[7] Analyzing dynamic discrete choice problems became popular with the seminal work of John Rust (Rust, 1987) who proposed an algorithm to estimate single agent dynamic models. Because his procedure is quite demanding when applied to a large state space, as in our experiment, an active literature has developed to improve the performance and feasibility of the estimator (e.g., Aguirregabiria and Mira, 2002; Arcidiacono and Miller, 2011), and to provide new identification results (e.g., Magnac and Thesmar, 2002; Blevins, 2014; Buchholz *et al.*, 2016), thereby increasing the opportunities for application in many areas of microeconomics (see Keane *et al.* (2011) for a recent survey in labor economics). Abbring (2010) and Aguirregabiria and Mira (2010) provide two technical surveys for the identification and estimation of dynamic discrete choice models.

than risk loving ones and, hence, the selling should happen earlier (as they do not need to wait longer for higher draws). It also should be noted that the predicted difference between the most risk loving and most risk averse participants is rather small, which suggests that any large differences in behavior cannot be attributed to heterogeneity in risk aversion.

Next we look at the relation between selling choices and the past peak. Overall, participants sell in 51% of the situations when the price goes above the past peak in the previous period. This provides evidence that the past peak has a strong influence on the decisions to sell (Prediction 1). The next question is What modulates the participants' decision to keep the asset despite reaching the past peak? We group new past peaks by how high they are and find that when the new past peak is above €8 selling happens in 71% of cases, in the range $[7, 8]$ – in 63%; in the range $[6, 7]$ – in 30%; and in the range $[5, 6]$ – in 2.6% (see Panel A in Figure 2). Notice that the percentages of selling, when the price is in the same intervals, but is *not* a new peak, are 46%, 32%, 14%, and 3.4% respectively, much lower values. This suggests that the past peak is an important reference point and that when participants decide to keep the asset they take into account the possibility that the price can grow higher in the future.

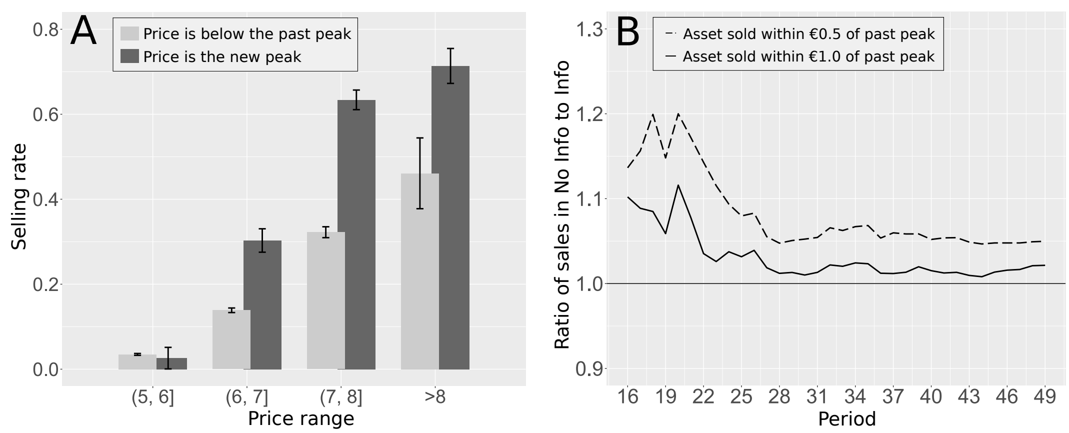

Figure 2: Panel A: The percentage of sales when the price reaches a new peak (dark grey) and when the price is below the current past peak (light grey). The error bars are the 95% confidence intervals. Panel B: The ratio of the number of sales up to period t in the No Info condition to the Info condition. The dashed line includes sales within €0.5 of the past peak and the solid line within €1 of the past peak.

Panel B in Figure 2 illustrates the difference in decisions to sell between the No Info and Info conditions. The two lines represent the cumulative ratio of the number of sales in the two conditions which are within €0.5 and €1 of the past peak. For each time period this ratio exceeds one which implies that there are more decisions to sell in the No Info than in the Info condition. This effect is especially evident in the early periods. In the late periods the number of selling decisions becomes approximately the same (Figure 13 in Appendix G shows that the ratios starting from period 33 oscillate close to one). This provides evidence, as we demonstrate with the structural model below, 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.[8] Moreover, the ratio is higher for the sales which are €0.5 close to the past peak than for the sales which are €1 close. 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 can be less salient due to the possibility of high prices after selling.

Pr[choice = keep]	I	II	III	IV	V
period	–0.101***	–0.104***	–0.106***	–0.100***	–0.098***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
price	0.175	0.116	0.228	0.151	0.113
	(0.137)	(0.135)	(0.132)	(0.136)	(0.138)
price^2	–0.170***	–0.166***	–0.177***	–0.173***	–0.170***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
entry price	0.470***	0.464***	0.450***	0.451***	0.452***
	(0.024)	(0.024)	(0.025)	(0.024)	(0.024)
future expected price		0.626*	0.649*	0.635**	0.609*
		(0.274)	(0.269)	(0.246)	(0.240)
past peak			0.216***	0.342***	0.280***
			(0.038)	(0.043)	(0.047)
future expected peak				0.276***	0.380***
				(0.045)	(0.057)
past peak×info				–0.283***	–0.153*
				(0.046)	(0.064)
info					–1.785**
					(0.639)
constant	8.044***	5.229***	3.358*	2.628*	3.280**
	(0.395)	(1.321)	(1.317)	(1.190)	(1.178)
N	112,137	112,137	112,137	112,137	112,137
BIC	31,272	31,268	31,236	31,209	31,213

Table 1: Random effects logit regression of the choice to keep the asset. 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). Errors are clustered by participant. The descriptions of all variables can be found in Appendix H.
***,** ,* denote statistical significance at the 0.1, 1 and 5 percent level.

To investigate further the influence of a larger set of independent variables on the choices to sell we run a logit regression presented in Table 1. The dependent variable choice is equal to 1 if the participant chooses to *keep the asset* in the current period and 0 if she decides to sell. First, we make several observations about the control variables. We see that the probability of selling increases with time (coefficient on period is negative). The negative coefficient on price^2 suggests non-linearity in response to price changes and increase in probability of selling as price increases. Variable entry price is equal to the price in period 15 at which the object was "bought" and is constant for all periods within a market. A positive coefficient could mean that the participants are loss averse and wait for the price to go above the entry price. However, the correlation between sale price and entry price is only 0.025, which does not support the loss aversion hypothesis. No-

[8] A regression table (Table 11) in Appendix M shows a similar result.

ticeably, the coefficient is always around 0.450 for all specifications that we report (Tables 5 and 6 in Appendix I). Thus, the loss aversion effect does not interact with other variables. Positive coefficient on future expected price shows that the selling behavior is modulated by future considerations, in particular, higher expected price in the future makes participants keep the asset longer. This is in line with the findings on off-peak choices to sell reported above. Finally, in Table 5 (Appendix I) we look at a regression that also includes personal measure of risk aversion from the Holt and Laury task. The coefficients on the Holt and Laury threshold (variable hl) are significant, negative, small, and do not affect the estimation of the main variables. The sign of the coefficients is in line with Prediction 4.

Now we discuss the main variables of interest: the market condition (info), the past peak (maximum price achieved in the past) and the future expected peak (expected future highest price). The coefficient on past peak is positive and very significant (column III), which shows that past regret indeed influences decisions to sell. The higher is the past peak the longer participants keep the asset waiting for the price to get closer to it. We account for the differences across conditions by adding the variables future expected peak, which is non-zero only in the Info condition, and the interaction of the past peak with info. The coefficients on both variables are highly significant (column IV). This points towards an important interaction between past and future regret in the Info condition: when future prices are observable the attention shifts from past regret to future regret thus diminishing the salience of the past peak. Overall, the regression in column IV directly shows the effect of information about the availability of future prices on the decisions to sell of our participants: just *knowing* that future prices will be observed after selling partially shifts the focus of the participants from the past prices to the future ones.[9]

It should be also mentioned that the regressions in columns III and IV can be interpreted in terms of a *distance* between the current price and the past and future expected peaks. All that is necessary is to split the coefficient on the variable price into several terms. For example, the regression in column III can be seen as the regression with coefficient 0.216^{***} on variable past peak – price and coefficient $0.443^{**} = 0.228 + 0.216^{***}$ on price.[10] Similar reinterpretation is possible for the regression in column IV. Therefore, in terms of the distances from the peaks, the regressions show that the probability of keeping the asset increases the farther away is the current price from the past and possibly future peak.

Finally, we test if the behavior of the participants changes as they gain experience during the experiment. We run the same regressions as in Table 1 only dividing markets into early (from 1 to 24) and late (from 25 to 48). Table 6 in Appendix I presents the results. The coefficients on the variables past peak, future expected peak and the interaction of past peak with info are significant,

[9] Column V shows the regression as in column IV, but with the variable Info added. It changes somewhat the coefficients on future expected peak and past peak×info and is itself significant. Nevertheless, we prefer to discuss the regression in column IV since it is stronger in terms of BIC.

[10] Alternatively, the variable distance, defined as past peak – price, could have been included directly in the regression. This, however, creates severe collinearity issues between this variable and price.

very similar for early and late markets, and very close to the coefficients in Table 1. This makes us conclude that there are no prominent learning effects and that the behavior of the participants does not change throughout the experiment.

6 Estimation of the Structural Model

We now turn to the estimation of the dynamic discrete choice model in Section 4. However, before proceeding to the estimation of (4.2) we analyze how the CCP differs in the two conditions, as this can further elucidate the mechanisms at play.

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 $t \in \{16, ..., 49\}$ 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. This is because the experiment has two conditions. The CCP for the No Info condition depends only on the price and the running past maximum of the process:

$$\Pr\{d = 0|\textsf{No Info}, x_t\} = \Lambda(\beta_{1t}\, y_t + \beta_{2t}\, s_{p,t})\ \forall t, \tag{6.1}$$

while in the Info condition it also depends on the expected future maximum:

$$\Pr\{d = 0|\textsf{Info}, x_t\} = \Lambda(\beta_{1t}\, y_t + \beta_{2t}\, s_{p,t} + \beta_{3t}\, s_{f,t})\ \forall t. \tag{6.2}$$

To maintain symmetry, the two logistic regressions are very similar. 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 may undermine the identification of the parameters.[11]

Figure 3 shows the projections of the time-averaged fitted CCP in the No Info condition. Specifically, each line represents the estimate of the probability of selling which results from averaging the fitted values of 34 logit regressions (one for each time period).[12] For the 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 zero 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

[11] Clustering at subject level does not affect the results. 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 matrix of covariates shows that including these terms makes it ill-conditioned in most periods.

[12] Thus, the CCPs in Figures 3 and 4 are shown just for illustration. They are out of sample estimates which 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).

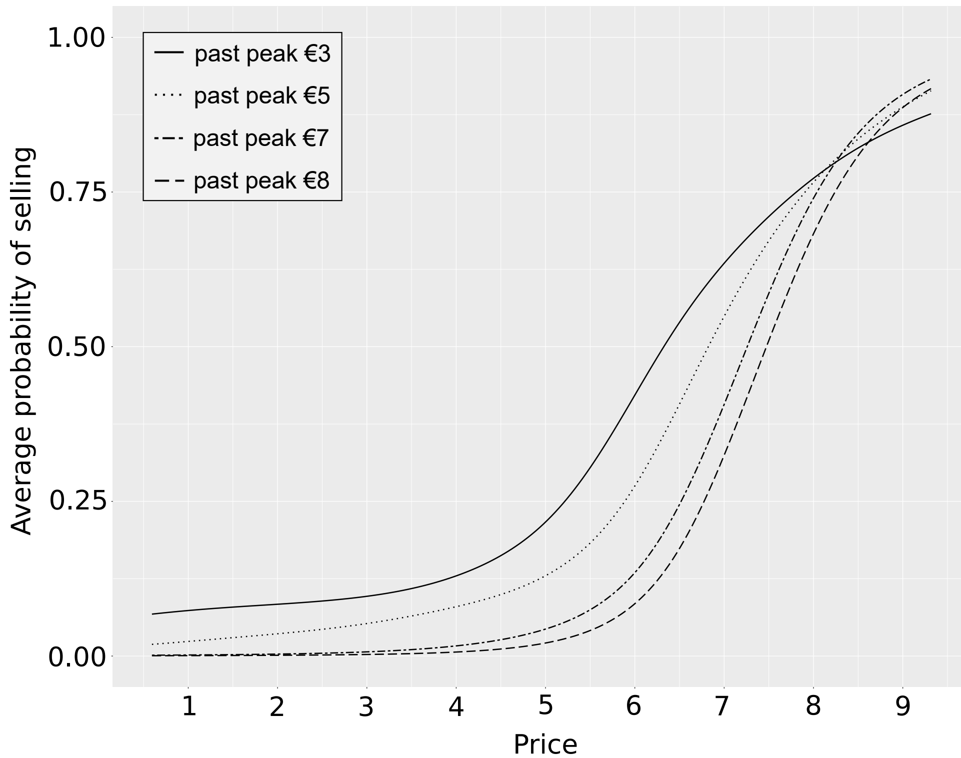

Figure 3: The effect of the past peak on the probability of selling the asset in the No Info condition. The conditional choice probability is computed by taking the average of the fitted values from (6.1) for all periods $t \in \{16, ..., 49\}$.

to become closer to it which corroborates Prediction 1. For the past peaks €7 and €8, which are very common in our data, the probability of selling increases rapidly when the price approaches €7. This strongly supports our theoretical predictions in Section 3 and demonstrates that the past peak indeed serves as a reference point.

Figure 4 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 3. 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 which dominates the past regret. In what follows we show that there is a *substitution effect* between the past and future regret that can explain this pattern.

In order to causally connect regret avoidance and decisions to sell in our experiment, we estimate (4.3) by non-linear least squares procedure. Nonparametric identification is shown in detail in Appendix F. We only provide an intuition of the proof here, which is standard. The most important step is to realize that the value function of the continuation choice (alternative 1) is a contraction mapping. Therefore there is a unique solution to $v^1(x_t)$. In addition, the difference of the two value functions is obtained using the formula for the CCP (Hotz and Miller, 1993). Given that selling is a terminating action ($v^0(x_t) = u^0(x_t)$), the per period utility, $u^0(x_t)$, is found by summing the CCP with $v^1(x_t)$.

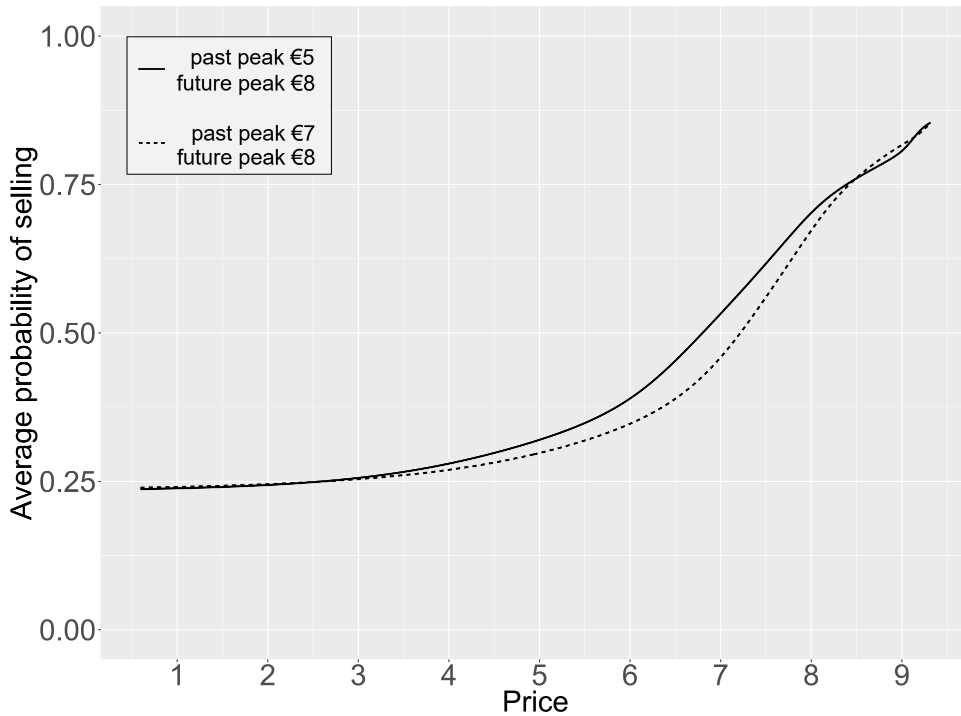

Figure 4: The effect of the past peak and the expected future peak in the Info condition. The conditional choice probability is computed by taking the average of the fitted values from (6.2) for all periods $t \in \{16, ..., 49\}$.

We estimate a parametric version of (4.3) which mirrors that provided in Section 3. The deterministic part of the per period utility is defined as follows:

$$u^0(y_t, s_{p,t}, s_{f,t}) = U(y_t; \rho) - R(s_{p,t}, s_{f,t}; \rho), \tag{6.3}$$

where $U(y_t; \rho) = \frac{y_t^{(1-\rho)} - 1}{1-\rho}$ is CRRA with risk aversion parameter $\rho \neq 1$ and $R(\cdot, \cdot; \rho)$ is the regret function and is defined as:

$$\begin{aligned} R(s_{p,t}, s_{f,t}; \rho) = \mathbb{1}_{\{\text{No Info}\}}\, \omega_{NI}\, U(s_{p,t}; \rho) \\ + \mathbb{1}_{\{\text{Info}\}}\, \big(\omega_I\, U(s_{p,t}; \rho) + \alpha_I\, U(s_{f,t}; \rho) + \lambda_I\, U(s_{p,t}; \rho) \times U(s_{f,t}; \rho)\big). \end{aligned} \tag{6.4}$$

The arguments of the regret function are the conditions (subscripts "NI" for No Info and "I" for Info), and the past and expected future peaks. The indicator function distinguishes the utility derived in one condition from another. An interesting feature of the regret function is the presence of an interaction term between past and future peaks for the Info condition. The interaction term captures the cross-partial derivative of the regret function while controlling for the risk aversion, which allows us to understand the degree of complementarity or substitutability of the two peaks. The parameters of $R(\cdot; \rho)$ are free to vary and indicate how strongly participants' decisions are affected by regret. Note, in fact, that, if ω_{NI}, ω_I, α_I, and λ_I are not significantly different from zero, the participants are categorized as regret neutral.

Table 2 displays the results of the estimation of (4.3) with the regret term (6.4) by nonlinear least squares on the dataset including periods $t \in \{16, ..., 48\}$.[13] All the results shown in the following tables hold also when the standard errors from the estimation of the CCP are clustered at subject level. The identification assumes that the discount factor is known, so the table shows utility function coefficients for $\beta \in \{0.99, 0.98, 0.97\}$. The results are robust across different designs, discount factors and discretization of the support.[14] Because of the nature and the length of the experiment, we expect β to be close to 1 and, therefore, the first column of Table 2 is our preferred model.

Parameter	$\beta = 0.99$	$\beta = 0.98$	$\beta = 0.97$
$\hat{\rho}$	-0.334^{***}	-0.335^{***}	-0.335^{***}
	(0.002)	(0.002)	(0.002)
$\hat{\omega}_{NI}$	0.988^{***}	1.006^{***}	1.002^{***}
	(0.068)	(0.034)	(0.022)
$\hat{\omega}_I$	1.200^{***}	1.072^{***}	1.023^{***}
	(0.099)	(0.064)	(0.054)
$\hat{\alpha}_I$	1.129^{***}	0.809^{***}	0.668^{***}
	(0.160)	(0.131)	(0.110)
$\hat{\lambda}_I$	-0.090^{***}	-0.063^{***}	-0.052^{***}
	(0.015)	(0.012)	(0.011)
N	111,613	111,613	111,613

Table 2: The estimation of (4.3) with the regret term (6.4) in periods 16 to 48 for different values of the discount factor β. Standard errors are in parenthesis.
***,** ,* denote statistical significance at the 0.1, 1 and 5 percent level.

Notice that $\hat{\omega}_{NI}$ and $\hat{\omega}_I$ are positive and significant. This supports the hypothesis that in both conditions participants experience disutility from past regret which is reflected in their selling choices. Most notably, the coefficient $\hat{\alpha}_I$ is also positive and significant which means that in the Info condition our participants are also influenced by future regret avoidance. Moreover, all these coefficients are large in magnitude as compared to the utility derived from selling the asset which has coefficient normalized to 1 (see equations 6.3 and 6.4). We also find that subjects are moderately risk seeking as $\hat{\rho}$ is negative but small.

The utility parameters estimates in Table 2 provide strong support for our hypotheses that past and future regret avoidance play a significant role in the decisions to sell the asset. However, the utility as expressed in (6.3) and (6.4) cannot tell us if future regret is *only* operational in the Info condition since this specification excludes any future influences in the No Info condition.

[13]For 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.

[14]Estimations for different regret functions $R(\cdot)$ and different specifications of the past maximum and discount factors are provided in Appendix J.1. The results with different discretization of the support (300 instead of 200 points) are reported in Appendix J.2.

Indeed, some evidence that future is taken into account in the No Info condition comes from the regression analysis in Section 5 where the variable future expected price significantly affects the probability to sell the asset (Table 1). In order to show that future regret is not playing a role in the No Info condition we estimate an extended structural model with past and future regret terms in both conditions. Table 3 shows the estimated parameters of the utility function with the regret term

$$\begin{aligned} R(s_{p,t}, s_{f,t}; \rho) = \mathbb{1}_{\{\text{No Info}\}} \left(\omega_{NI} U(s_{p,t}; \rho) + \alpha_{NI} U(s_{f,t}; \rho) + \lambda_{NI} U(s_{p,t}; \rho) \times U(s_{p,t}; \rho) \right) \\ + \mathbb{1}_{\{\text{Info}\}} \left(\omega_I U(s_{p,t}; \rho) + \alpha_I U(s_{f,t}; \rho) + \lambda_I U(s_{p,t}; \rho) \times U(s_{f,t}; \rho) \right). \end{aligned} \quad (6.5)$$

Overall, the parameter estimates of past and future regret in the Info condition are the same as in Table 2. 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.

Parameter	$\beta = 0.99$	$\beta = 0.98$	$\beta = 0.97$
$\hat{\rho}$	-0.322^{***} (0.003)	-0.321^{***} (0.003)	-0.321^{***} (0.003)
$\hat{\omega}_{NI}$	0.640^{***} (0.120)	0.800^{***} (0.077)	0.857^{***} (0.065)
$\hat{\omega}_I$	1.242^{***} (0.124)	1.133^{***} (0.081)	1.100^{***} (0.068)
$\hat{\alpha}_{NI}$	0.155 (0.190)	0.157 (0.156)	0.183 (0.131)
$\hat{\alpha}_I$	1.000^{***} (0.201)	0.706^{***} (0.166)	0.587^{***} (0.139)
$\hat{\lambda}_{NI}$	0.000 (0.018)	-0.003 (0.015)	-0.008 (0.013)
$\hat{\lambda}_I$	-0.084^{***} (0.019)	-0.060^{***} (0.016)	-0.052^{***} (0.014)
N	111,613	111,613	111,613

Table 3: The estimation of (4.3) with the regret term (6.5) in periods 16 to 48 for different values of the discount factor β. Standard errors are in parenthesis. The CCP is computed using the formula in (6.2) for both conditions.
***,**,* denote statistical significance at the 0.1, 1 and 5 percent level.

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. This can be interpreted as 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 among 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 5 in Section 7). Moreover, the presence of the interaction term implies that participants switch their focus between the 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.

Finally, we verify that our results cannot be explained by heterogeneity in risk aversion, as participants with different risk preferences would sell at different times (though, in any way, risk preferences cannot account for the difference between the No Info and Info conditions). In Section 5 we have already provided evidence that different risk preferences do not affect the estimates for past and future regret and have small overall impact on the decisions to sell. In Appendix J.3 we report the estimation of a model where agents are assumed to belong to two types, different in their risk preferences. The estimation confirms our results in Table 3 in the sense that the estimates of past and future regret stay unchanged while both types are mildly risk seeking.

7 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 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. This means that the participants try to minimize the distance of the selling price to the highest peak be it in the past or in the future.

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 act 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*. Indeed, the regression analysis in Section 5 shows that in the No Info condition the future expected price of the asset does have a say in the decision to sell the asset which means that the participants do think about the future, but just choose to ignore the prospect of feeling 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.[15]

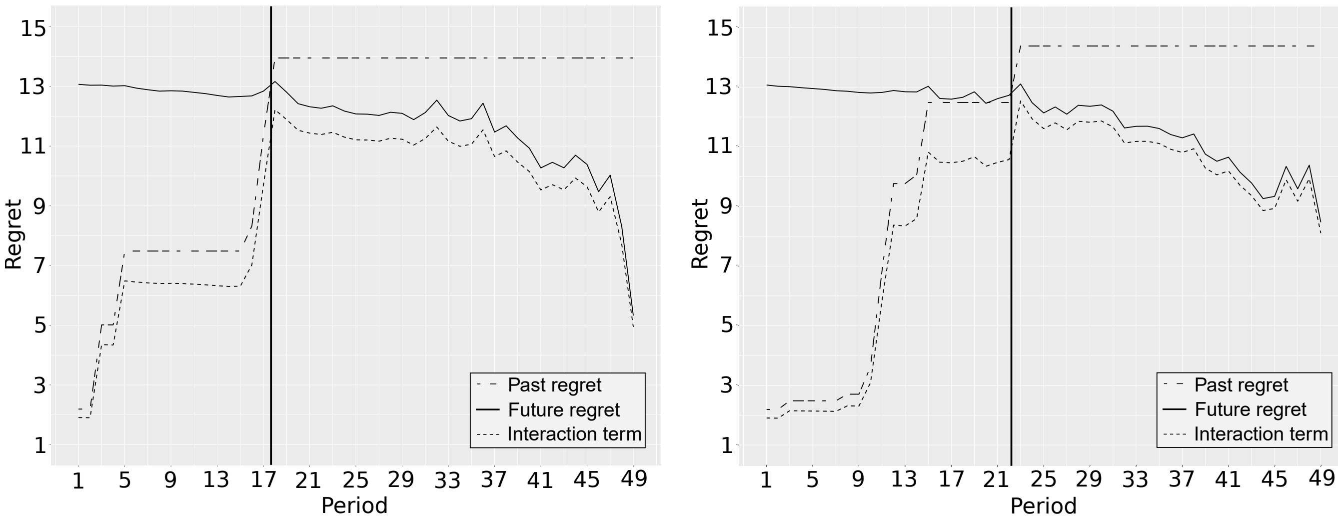

Figure 5: Examples of the dynamics of past and future regret in two selected markets in the Info condition for the periods $t \in \{16, ..., 48\}$. The curves show the terms of the estimated regret function (i.e., past regret = $\hat{\omega}_I \, U(s_{p,t}; \rho)$, future regret = $\hat{\alpha}_I \, U(s_{f,t}; \rho)$ and interaction term = $|\hat{\lambda}_I| \, U(s_{f,t}; \rho) \, U(s_{p,t}; \rho)$). The solid vertical line shows the moment at which the participants switch the focus from future regret to past regret.

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 5 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 which 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 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 2. In the early periods future regret is a reference point for the participants in the Info condition but not for the participants in the No

[15] The 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.

Info condition. As time goes by, the saliency of past regret increases eventually dominating the future regret term (see Figure 5 and Figure 14 in Appendix G). This makes behavior in the two conditions observationally equivalent.

In our experiment this effect is detected *within subjects*, which means that orientation towards 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 *et al.*, 1998; Bordalo *et al.*, 2017). 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 is the end, the lower is 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.

Our results imply another interesting behavioral effect which is concerned with the potential choice between observing and not observing the future price after selling the market. 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 (Strahilevitz *et al.*, 2011). This can have consequences for policies directed at 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 which 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.

8 Conclusion

In an experimental task which 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 which they would experience if the price of the asset increased after they sold it and 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 which 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 policy makers.

References

ABBRING, J. H. (2010). Identification of dynamic discrete choice models. *Annual Review of Economics*, **2** (1), 367–394.

AGUIRREGABIRIA, V. and MIRA, P. (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica: Journal of the Econometric Society*, **70** (4), 1519–1543.

— and — (2010). Dynamic discrete choice structural models: a survey. *Journal of Econometrics*, **156** (1), 38–67.

ARCIDIACONO, P. and MILLER, R. A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica: Journal of the Econometric Society*, **79**, 1823–1867.

BAJARI, P. and HORTAÇSU, A. (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy*, **113** (4), 703–741.

BAUCELLS, M., WEBER, M. and WELFENS, F. (2011). Reference-point formation and updating. *Management Science*, **57** (3), 506–519.

BEBER, A. and PAGANO, M. (2013). Shortselling bans around the world: evidence from the 2007-09 crisis. *The Journal of Finance*, **68** (1), 343–381.

BELL, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, **30** (5), 961–981.

BIKHCHANDANI, S. and SEGAL, U. (2014). Transitive regret over statistically independent lotteries. *Journal of Economic Theory*, **152**, 237–248.

BLEICHRODT, H., CILLO, A. and DIECIDUE, E. (2010). A quantitative measurement of regret theory. *Management Science*, **56** (1), 161–175.

BLEVINS, J. R. (2014). Nonparametric identification of dynamic decision processes with discrete and continuous choices. *Quantitative Economics*, **5** (3), 531–554.

BOERI, M., LONGO, A., GRISOLA, J. M., HUTCHINSON, W. G. and KEE, F. (2013). The role of regret minimisation in lifestyle choices affecting the risk of coronary heart disease. *Journal of Health Economics*, **32** (2013), 253–260.

BORDALO, P., GENNAIOLI, N. and SHLEIFER, A. (2017). Diagnostic expectations and credit cycles. *Journal of Finance*, **forthcoming**.

BUCHHOLZ, N., SHUM, M. and XU, H. (2016). Semiparametric estimation of dynamic discrete choice models, Mimeo, Princeton University, Caltech and University of Texas at Austin.

CAMILLE, N., CORICELLI, G., SALLET, J., PRADAT-DIEHL, P., DUHAMEL, J.-R. and SIRIGU, A. (2004). The involvement of the orbitofrontal cortex in the experience of regret. *Science*, **304** (5674), 1167–1170.

CORICELLI, G., CRITCHLEY, H. D., JOFFILY, M., O'DOHERTY, J. P., SIRIGU, A. and DOLAN, R. J. (2005). Regret and its avoidance: a neuroimaging study of choice behavior. *Nature Neuroscience*, **8** (9), 1255–1262.

CRAWFORD, V. P. and IRIBERRI, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica: Journal of the Econometric Society*, **75** (6), 1721–1770.

ESPONDA, I. and VESPA, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, **6** (4), 180–202.

FILIZ-OZBAY, E. and OZBAY, E. Y. (2007). Auctions with anticipated regret: theory and experiment. *The American Economic Review*, **97** (4), 1407–1418.

FISCHBACHER, U. (2007). z-tree: Zurich Toolbox for ready-made economic experiments. *Experimental Economics*, **10** (2), 171–178.

FOGEL, S. O. and BERRY, T. (2006). The disposition effect and individual investor decisions: the roles of regret and counterfactual alternatives. *The Journal of Behavioral Finance*, **7** (2), 107–116.

FRYDMAN, C. and CAMERER, C. (2016). Neural evidence of regret and its implications for investor behavior. *Review of Financial Studies*, **forthcoming**.

—, HARTZMARK, S. M. and SOLOMON, D. H. (2017). Rolling mental accounts. *Review of Financial Studies*, **forthcoming**.

GILLEN, B. J. (2009). Identification and estimation of level-k auctions, Mimeo, Caltech.

GNEEZY, U. (2005). *Updating the reference level: experimental evidence*, Springer US, chap. 12, pp. 263–284.

HAYASHI, T. (2008). Regret aversion and opportunity dependence. *Journal of Economic Theory*, **139** (1), 242–268.

— and YOSHIMOTO, H. (2016). Risk-and regret-averse bidders in sealed-bid auctions. *Working Paper SSRN 2400092*.

HAZAN, E. and KALE, S. (2015). An online portfolio selection algorithm with regret logarithmic in price variation. *Mathematical Finance*, **25** (2), 288–310.

HOLT, C. A. and LAURY, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, **92** (5), 1644–1655.

HOTZ, V. J. and MILLER, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, **60** (3), 497–529.

KEANE, M. P., TODD, P. E. and WOLPIN, K. I. (2011). The structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. *Handbook of Labor Economics*, **4**, 331–461.

KLIBANOFF, P., LAMONT, O. and WIZMAN, T. A. (1998). Investor reaction to salient news in closed-end country funds. *The Journal of Finance*, **53** (2), 673–699.

LEUNG, S. and HALPERN, J. Y. (2015). Minimizing regret in dynamic decision problems. *Working Paper available at arXiv 1502.00152*.

LOOMES, G. and SUGDEN, R. (1982). Regret theory: an alternative theory of rational choice under uncertainty. *The Economic Journal*, **92** (368), 805–824.

MAGNAC, T. and THESMAR, D. (2002). Identifying dynamic discrete decision processes. *Econometrica: Journal of the Econometric Society*, **70** (2), 801–816.

MCFADDEN, D. L. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (ed.), *Frontiers in Econometrics*, Academic Press, New York, pp. 105–142.

MICHENAUD, S. and SOLNIK, B. (2008). Applying regret theory to investment choices: currency hedging decisions. *Journal of International Money and Finance*, **27** (5), 677–694.

MUERMANN, A., MITCHELL, O. S. and VOLKMAN, J. M. (2006). Regret, portfolio choice, and guarantees in defined contribution schemes. *Insurance: Mathematics and Economics*, **39** (2), 219–229.

OPREA, R. (2014). Survival versus profit maximization in a dynamic stochastic experiment. *Econometrica: Journal of the Econometric Society*, **82** (6), 2225–2255.

—, FRIEDMAN, D. and ANDERSON, S. T. (2009). Learning to wait: a laboratory investigation. *The Review of Economic Studies*, **76** (3), 1103–1124.

QIN, J. (2015). A model of regret, investor behavior, and market turbulence. *Journal of Economic Theory*, **160**, 150–174.

RUST, J. (1987). Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society*, **55** (5), 999–1033.

— (1994). *Handbook of Econometrics*, Amsterdam. North Holland, vol. 4, chap. Structural estimation of Markov decision processes, pp. 3081–3143.

SARVER, T. (2008). Anticipating regret: why fewer options may be better. *Econometrica: Journal of the Econometric Society*, **76** (2), 263–305.

SAVAGE, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons, Inc., New York.

Skiadas, C. (1997). Conditioning and aggregation of preferences. *Econometrica: Journal of the Econometric Society*, **65** (2), 347–367.

Strack, P. and Viefers, P. (2015). Too proud to stop: regret in dynamic decisions. *Working Paper SSRN 2465840*.

Strahilevitz, M. A., Odean, T. and Barber, B. M. (2011). Once burned, twice shy: how naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research*, **48** (SPL), S102–S120.

Tauchen, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, **20**, 177–181.

Zeelenberg, M. and Beattie, J. (1997). Consequences of regret aversion 2: additional evidence for effects of feedback on decision making. *Organizational Behavior and Human Decision Processes*, **72** (1), 1997.

—, —, van der Pligt, J. and de Vries, N. K. (1996). Consequences of regret aversion: effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes*, **65** (2), 148–158.

Supplementary Material (for Online Publication)

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 K). 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 6 and 7 explain.

The timing of each market was as following. The new price was shown each 0.8 seconds.[1] This was

[1]The 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.

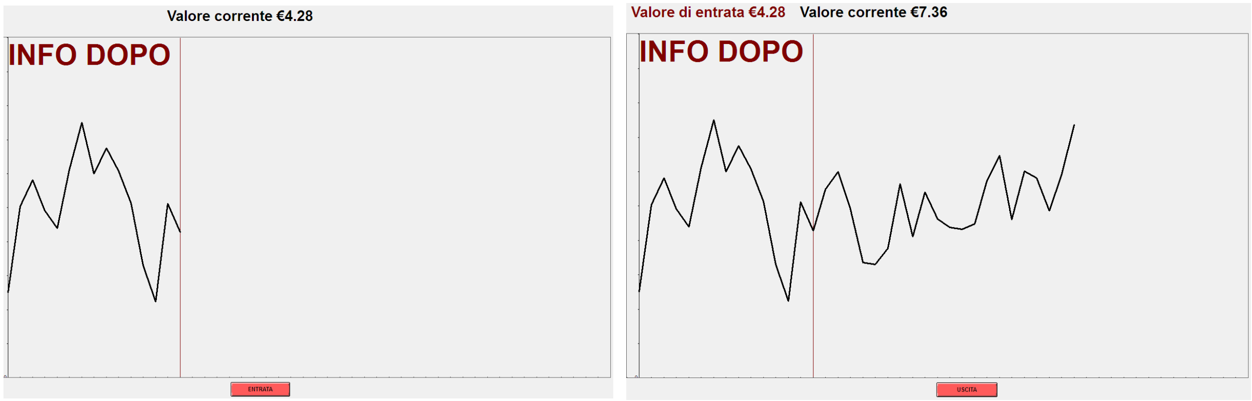

Figure 6: The left picture shows the market price evolution before period 15, which is marked by a vertical red line. At period 15 the market "stopped," so that participants could inspect the entry price. An ENTER (ENTRATA) button should have been pressed to start the market again. After period 15 the 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 7: 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."

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.

A.2 Information about the Price Dynamics and Training

Participants were explicitly informed about the process that generates the price (see instructions in Appendix K). 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 8 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 7). 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 Halt 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 and Vostroknutov (2016). The instructions and the screenshots are presented in Appendix L.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

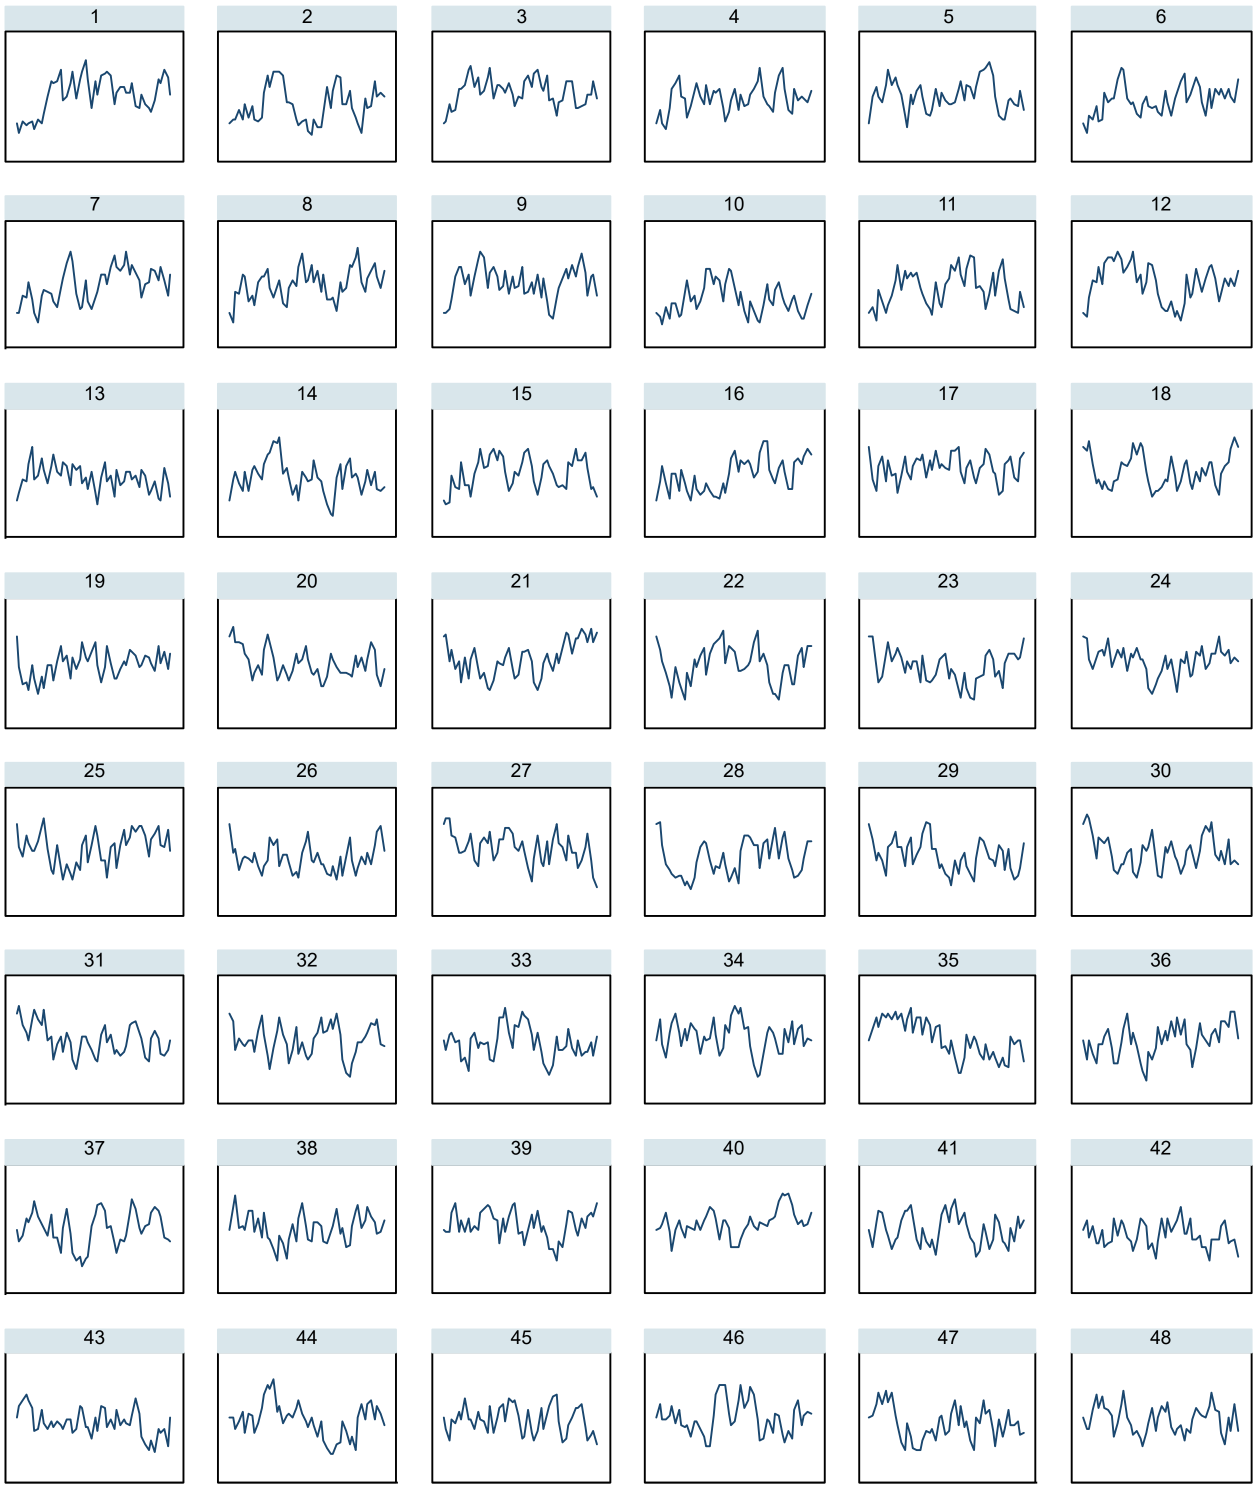

Figure 8: Prices in 48 markets.

B Proof of Prediction 4

Proposition 1. *In the case without regret risk averse agents sells earlier than risk seeking agents.*

Proof. An agent without regret sells if

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

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

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

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:

$$\tilde{y}_t = \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\tilde{v}_{t+2}|\tilde{y}_t]\} \tag{B.3}$$

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*

$$\tilde{y}_t^{1-\rho} \geq \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|\tilde{y}_t]\}. \tag{B.4}$$

Plugging (B.3) into (B.4) we get

$$\max\{\mathbb{E}_{y_{t+1}}[y_{t+1}|\tilde{y}_t]^{1-\rho}, \mathbb{E}_{y_{t+1}}[\tilde{v}_{t+2}|\tilde{y}_t]^{1-\rho}\} \geq \max\{\mathbb{E}_{y_{t+1}}[y_{t+1}^{1-\rho}|\tilde{y}_t], \mathbb{E}_{y_{t+1}}[\dot{v}_{t+2}|\tilde{y}_t]\}. \tag{B.5}$$

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

$$\begin{gathered}\mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho} = \left(\sum_\iota \Pr\{y_{T-1,\iota}|y_{T-2}\}\mathbb{E}_{y_T}[y_T|y_{T-1,\iota}]\right)^{1-\rho} \\ \text{and} \\ \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}] = \sum_\iota \Pr\{y_{T-1,\iota}|y_{T-2}\}\mathbb{E}_{y_T}[y_T^{1-\rho}|y_{T-1,\iota}]\end{gathered} \tag{B.6}$$

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.6) can be rewritten as

$$\begin{gathered}\left(\sum_\iota \Pr\{y_{T-1,\iota}|y_{T-2}\}\sum_{\xi_\iota}\Pr\{y_{T,\xi_\iota}|y_{T-1,\iota}\}y_{T,\xi_\iota}\right)^{1-\rho} = \left(\sum_\zeta p_\zeta y_{T,\zeta}\right)^{1-\rho} \\ \text{and} \\ \sum_\iota \Pr\{y_{T-1,\iota}|y_{T-2}\}\sum_{\xi_\iota}\Pr\{y_{T,\xi_\iota}|y_{T-1,\iota}\}y_{T,\xi_\iota}^{1-\rho} = \sum_\zeta p_\zeta y_{T,\zeta}^{1-\rho}\end{gathered} \tag{B.7}$$

respectively. Here ξ_ι enumerates y_T for each ι and ζ enumerates all combinations of ι and ξ_ι. Now, the RHS of the first equation in (B.7) 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

$$\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}|y_{T-2}]^{1-\rho}, \mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho}\} > \max\{\mathbb{E}_{y_{T-1}}[y_{T-1}^{1-\rho}|y_{T-2}], \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]\}. \tag{B.8}$$

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.8) holds.

Since (B.8) holds for all y_{T-2}, it is true that

$$\begin{aligned}\mathbb{E}_{y_{T-2}}[\tilde{v}_{T-1}|y_{T-3}]^{1-\rho} = \mathbb{E}_{y_{T-2}}[\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}|y_{T-2}]^{1-\rho}, \mathbb{E}_{y_{T-1}}[\tilde{v}_T|y_{T-2}]^{1-\rho}\}|y_{T-3}] > \\ \mathbb{E}_{y_{T-2}}[\max\{\mathbb{E}_{y_{T-1}}[y_{T-1}^{1-\rho}|y_{T-2}], \mathbb{E}_{y_{T-1}}[\dot{v}_T|y_{T-2}]\}|y_{T-3}] = \mathbb{E}_{y_{T-2}}[\dot{v}_{T-1}|y_{T-3}].\end{aligned} \tag{B.9}$$

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.4) holds with strict inequality for all t as long as $\rho \in (0,1)$. When the agent is risk seeking, or $\rho < 0$, (B.4) 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.[2] 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

$$\mathbb{E}_{y_{t+1}}[...\mathbb{E}_{y_\tau}[y_\tau^{1-\rho}|y_{\tau-1}]...|y_t] = \mathbb{E}_{\varepsilon_{t+1}}...\mathbb{E}_{\varepsilon_\tau}[(\alpha^{\tau-t}y_t + (1-\alpha^{\tau-t})E_\tau)^{1-\rho}] \tag{B.10}$$

where E_τ is a weighted average of random variables $\varepsilon_{t+1}, \varepsilon_{t+2}, ..., \varepsilon_\tau$. All functions in (B.10) 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 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.10). Thus,

$$y^{1-\rho} = \sum_\iota p_\iota \mathbb{E}_{\varepsilon_{t+1}}...\mathbb{E}_{\varepsilon_{\tau_\iota}}[(\alpha^{\tau_\iota - t}y + (1-\alpha^{\tau_\iota - t})E_{\tau_\iota})^{1-\rho}] = \sum_\iota p_\iota \mathbb{E}_{\tau_\iota}[(\alpha^{\tau_\iota - t}y + (1-\alpha^{\tau_\iota - t})E_{\tau_\iota})^{1-\rho}] \tag{B.11}$$

for some enumeration $\{p_\iota, \tau_\iota\}_\iota$ and with $\mathbb{E}_{\tau_\iota}$ being short for $\mathbb{E}_{\varepsilon_{t+1}}...\mathbb{E}_{\varepsilon_{\tau_\iota}}$. Notice that the derivatives of functions (B.10) 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 $\mathbb{E}_\tau$ 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 summation. Keeping this in mind let us rewrite (B.11) as

$$(1-\rho)y^{-\rho} = \sum_\iota p_\iota \alpha^{\tau_\iota - t}(1-\rho)\mathbb{E}_{\tau_\iota}[(\cdot)^{-\rho}] + \frac{1-\rho}{y}\sum_\iota p_\iota \mathbb{E}_{\tau_\iota}[(1-\alpha^{\tau_\iota - t})E_{\tau_\iota}(\cdot)^{-\rho}] \tag{B.12}$$

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

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

$$(1-\rho)y^{-\rho} = \frac{dm(y)}{dy} + \frac{1-\rho}{y} \sum_\iota p_\iota \mathbb{E}_{\tau_\iota}[(1-\alpha^{\tau_\iota - t})E_{\tau_\iota}(\cdot)^{-\rho}]. \tag{B.13}$$

Here LHS is the derivative of LHS of (B.11) at y and RHS is the derivative of m at y plus a positive number. Thus, at y the derivative of $y_t^{1-\rho}$ 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 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 averse agent sells at a price where risk neutral agent is indifferent and that risk seeking agent keeps the asset at that price, we can conclude that risk averse agent must have selling threshold at a price below risk neutral agent and risk seeking agent must have the threshold above it. Therefore, risk averse agent, given the same prices, sells before risk neutral agent and risk seeking agent sells after. ■

C 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 $t \in \{2, ..., T\}$, $y_{t+1} = \alpha y_t + (1-\alpha)u_t$ 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 $k \in \{1, ..., T-2\}$, 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:

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

The distribution of y_t^k is

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

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:

$$\text{Future peak}_{\text{period k}} = \int_0^{10} v \mathrm{d} \prod_{j=1}^{T-k} F_{(j)}(v)$$
$$= \int_0^{10} v \sum_{j=1}^{T-k} f_{(j)}(v) \prod_{h \neq j}^{T-k} F_{(h)}(v) \mathrm{d}v$$
$$= \int_0^{10} v \sum_{j=1}^{T-k} f_{(j)}(v) \prod_{h \neq j}^{T-k} \int_0^v f_{(h)}(s) \mathrm{d}s \, \mathrm{d}v$$

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 (C.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 $[\alpha^{t-1} \frac{y_1}{t-1}, \alpha^{t-1} \frac{y_1}{t-1} + 10(1-\alpha)\alpha^{j-1}], \forall j \in \{1, ..., t-1\}$. According to Potuschak and Muller (2009, section 2.2.2, page 180), the density is

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

where . indicates the dot product, lower bar means vector, $a = \{5(1-\alpha), 5\alpha(1-\alpha), 5\alpha^2(1-\alpha), ...5\alpha^{t-1}(1-\alpha)\}$, $c = \{\alpha^{t-1}\frac{y_1}{t-1} + 5(1-\alpha), \alpha^{t-1}\frac{y_1}{t-1} + 5\alpha(1-\alpha), \alpha^{t-1}\frac{y_1}{t-1} + 5\alpha^2(1-\alpha), ...\alpha^{t-1}\frac{y_1}{t-1} + 5\alpha^{t-1}(1-\alpha)\}, \forall\, 1 \leq j \leq t-1$. σ_j and $\underline{\varepsilon}_j$ are matrices which deal with positive and negative signs (see Potuschak and Muller (2009) for more information). 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 $F_{(1)}(s) = \frac{s-\alpha y_1}{(1-\alpha)10}$.

C.1 Normal Approximation

(C.2) is problematic, because, as the number of uniform RVs to be summed increases, the denominator goes to zero since $a_k \to 0$. 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:

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

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:

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

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 $E[|X_i|^3] = \int_{-a_i}^{a_i} |s|^3 f_i(s)\mathrm{d}s = \int_{-a_i}^{a_i} |s|^3 \frac{1}{2a_i}\mathrm{d}s$. Solving the integral we get:

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

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 $\big(\frac{25}{3}\big)^{\frac{3}{2}}\big(\sum_i^N (1-\alpha)^2\alpha^{2(i-1)}\big)^{\frac{3}{2}}$ (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 quantites, the result is $W \times \frac{\sum_i^N (1-\alpha)^3\alpha^{3(i-1)}}{\big(\sum_i^N (1-\alpha)^2\alpha^{2(i-1)}\big)^{\frac{3}{2}}}$, where $0 < W < 1$ is a constant. Finally, we can establish that:

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

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

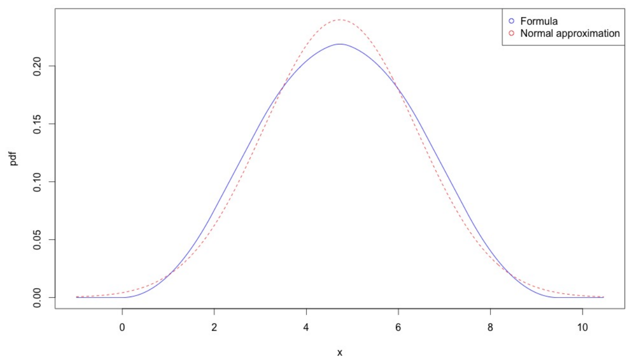

Figure 9: pdf, sum of 3 uniform RVs

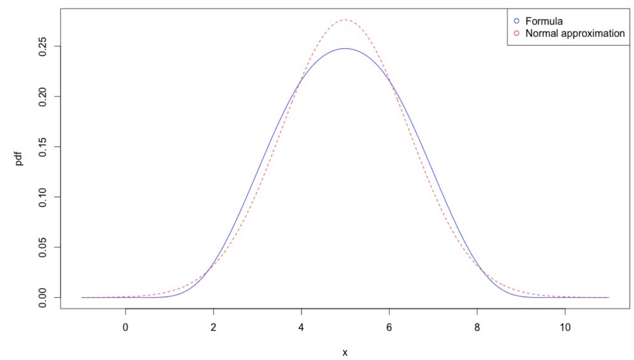

Figure 10: pdf, sum of 13 uniform RVs

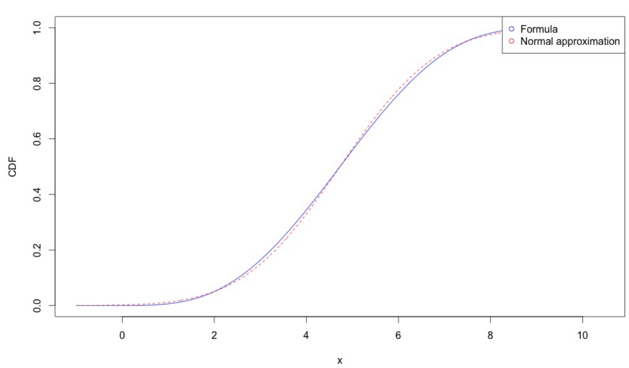

Figure 11: CDF, sum of 3 uniform RVs

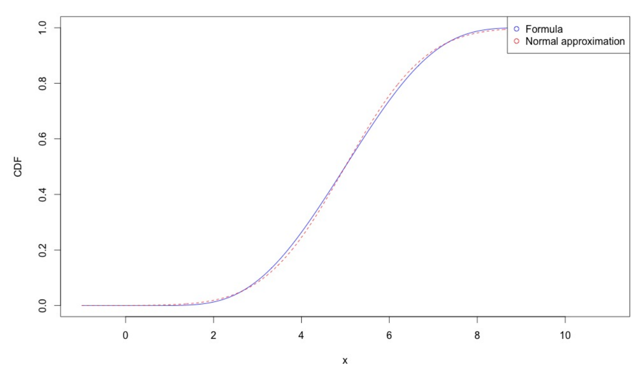

Figure 12: CDF, sum of 13 uniform RVs

D 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:

$$y_{i,t+1} = \mu + \rho y_{i,t} + \varepsilon \tag{D.1}$$

where $y_{i,t+1}, y_{i,t}$ are the prices for participant $i = \{1, ..., N\}$ at time $t+1$ and t respectively, and $\varepsilon \sim N(0, \sigma_i^2)$. 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 $\hat{\mu} = 1.9745$, $\hat{\rho} = 0.6015$ and $\hat{\sigma} = 1.1613$. The estimate of ρ seems to be very close to the parameter α which updates the price from period y_t to y_{t+1} ($\alpha = 0.6$).

Let $\{y^1, ..., y^K\}$ denote the support of the discretized variable $\tilde{Y}_{i,t}$, where $y^1 > y^2 > ... > y^{K-1} > y^K$ with $K = 200$ 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 €0.5937, upper bound (y^{200}) equal to €9.3163, and interval between adjacent points (ω) equal to €0.0438.

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,[3] 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 (D.1).

[3] The standardization implies that the distribution is a standard normal.

E 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. 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

$$\mathbb{E}\left\{\sum_{t=1}^{T}\beta^{t-1}u^{d}(x_t)\right\}$$

where $\beta \in (0,1)$ is a discount factor. 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.[4] d is the participant's binary choice at time $t \leq T$:

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

$u^d(x_t)$ is the payoff after choosing alternative d; 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 CRRA consumption utility function and $R(x_t)$ measures regret.

The flow (per period) payoff from choice d at period t is $u^d + \varepsilon_t^d$ where the error term ε_t^d is independent of x. As in Murphy (2015), the error term is assumed to be $\varepsilon^d = \tilde{\varepsilon}^d - \sigma_\varepsilon\gamma$ where $\tilde{\varepsilon}^d$ is distributed Type I extreme value with location parameter equal to zero and scale parameter $\sigma_\varepsilon = 1$[5]. By the properties of the Type I extreme value distribution, the mean of $\tilde{\varepsilon}^d$ is γ (the Euler's constant). ε^d is therefore mean zero. Given these preliminaries, denote by $V(x_t, \varepsilon_t) = \max_{d\in\{0,1\}}\{v^d(x_t) + \varepsilon^d\}$ the value function at the beginning of period t with $\varepsilon_t = \{\varepsilon_t^0, \varepsilon_t^1\}$ and define the *alternative specific value function* (ASVF) for option $d \in \{0,1\}$ at time t as:

$$v^d(x_t) = \begin{cases} 0 + \beta\mathbb{E}\{v(x_{t+1})|x_t, d=1\} & \text{if } d = 1 \text{ (keep)} \\ u^0(x_t) & \text{if } d = 0 \text{ (sell)} \end{cases} \tag{E.1}$$

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 (E.1), can be rewritten as the expectation over the error term, ε_t, of the value function at time t

$$v(x_t) \equiv \int V(x_t, \varepsilon_t)\mathrm{d}\Lambda(\varepsilon_t)$$

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

$$v^d(x_t) = u^d(x_t) + \beta\mathbb{E}\{v_{t+1}(x_{t+1})|x_t,\}, \;\; d \in \{0,1\}. \tag{E.2}$$

[4] Identification 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.

[5] The 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:

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

where $v^d(\cdot)$ is defined as in (E.2). Denote the *Conditional Choice Probability (CCP)* of continuing (action 1) as $\Pr\{d = 1|x_t\} \equiv p^1(x_t)$:

$$p^1(x_t) = \frac{\exp(v^1(x_t))}{\exp(v^1(x_t)) + \exp(v^0(x_t))} = \frac{1}{1 + \exp(v^0(x_t) - v^1(x_t))}. \tag{E.3}$$

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

$$\phi\big(p^1(x_t)\big) \equiv \ln\big(p^1(x_t)\big) - \ln\big(1 - p^1(x_t)\big) \equiv v^1(x_t) - v^0(x_t). \tag{E.4}$$

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

$$\begin{aligned} v^0(x_t) &= u^0(x_t) \\ v^1(x_t) &= 0 + \beta \int_{\mathcal{X}_{t+1}} \int_{\varepsilon} \max\{v^0(x_{t+1}) + \varepsilon_{t+1}^0, v^1(x_{t+1}) + \varepsilon_{t+1}^1)\} \mathrm{d}\Lambda(\varepsilon) \mathrm{d}F(x_{t+1}|x_t), \end{aligned} \tag{E.5}$$

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. $F(x_{t+1}|x_t, d = 0) = 0$). The estimation is based on the difference of the two ASVFs in (E.5):

$$v^1(x_t) - v^0(x_t) = -u(x_t) + \beta \int_{\mathcal{X}_{t+1}} \int_{\varepsilon} \max\{v^0(x_{t+1}) + \varepsilon_{t+1}^0, v^1(x_{t+1}) + \varepsilon_{t+1}^1)\} \mathrm{d}\Lambda(\varepsilon) \mathrm{d}F(x_{t+1}|x_t) \tag{E.6}$$

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

$$\begin{aligned} v^1(x_t) &= \beta \int_{\mathcal{X}_{t+1}} \int_{\varepsilon} \max\{v^0(x_{t+1}) + \varepsilon_{t+1}^0, v^1(x_{t+1}) + \varepsilon_{t+1}^1\} \mathrm{d}\Lambda(\varepsilon) \mathrm{d}F(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}_{t+1}} \gamma + \log\big(\exp(v^0(x_{t+1}) - \gamma) + \exp(v^1(x_{t+1}) - \gamma)\big) \mathrm{d}F(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}_{t+1}} \gamma + \log\Big(\big(1 + \exp(v^1(x_{t+1}) - v^0(x_{t+1}))\big) \exp(v^0(x_{t+1}) - \gamma)\Big) \mathrm{d}F(x_{t+1}|x_t) \\ &= \beta \int_{\mathcal{X}_{t+1}} \big(u^0(x_{t+1}) - \log(Pr\{d_{t+1} = 0|x_{t+1}\})\big) \mathrm{d}F(x_{t+1}|x_t) \end{aligned}$$

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

$$v^1(x_t) - v^0(x_t) = -u^0(x_t) + \beta \int_{\mathcal{X}_{t+1}} \big(v^0(x_{t+1}) - \log(Pr\{d_{t+1} = 0|x_{t+1}\})\big) \mathrm{d}F(x_{t+1}|x_t).$$

By replacing the dependent variable in the last equation with $\phi(p^1(x_t))$ and by discretizing the state space

$\mathcal{X}_t$ the objective function can be rewritten in an estimable form:

$$\begin{aligned}\phi(p^1(x_t)) &= -u^0(x_t) + \beta \sum_{\mathcal{X}_{t+1}} \left(v^0(x_{t+1}) - \log(Pr\{d_{t+1} = 0|x_{t+1}\})\right) f(x_{t+1}|x_t) \\ &= -u^0(x_t) + \beta \sum_{\mathcal{X}_{t+1}} \left(u^0(x_{t+1}) - \log(p^0(x_{t+1}))\right) f(x_{t+1}|x_t)\end{aligned}$$

which concludes the derivation.

Note that the regret components are functions of price (y_t is the only random variable) and time. In fact, $s_{p,t} = \max_{s \leq t} y_s$ 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.[6] Therefore, $Pr\{y_{t+1}, s_{p,t+1}, s_{f,t+1}|x_t, d = 1\} = f(y_{t+1}, s_{p,t+1}, s_{f,t+1}|y_t, s_{p,t}, s_{f,t}) = f(y_{t+1}, s_{p,t+1}, s_{f,t+1}|y_t, s_{p,t})$. 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

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

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

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

F Nonparametric Identification

Identification of the objects of interest is standard and the proof is exposed here for completeness.

The first step of the identification procedure requires the identification of the transition matrix of the Markovian process and of the conditional choice probabilities (CCP) (see equations 6.1 and 6.2), which are obtained directly from the data.

The second step involves the nonparametric identification of the utility function, and is standard. Consider the following assumptions:

Assumption 1: *Additive separability*. The flow utility is separable in the observables and unobservable arguments, $U(d, x_t, \varepsilon) = u^d(x_t) + \varepsilon^d$.

Assumption 2: *The unobservables are iid*. The unobservable state variables, $\varepsilon_t = (\varepsilon_t^0, \varepsilon_t^1)$, are iid across time. Moreover ε^d is distributed Type I extreme value.

Assumption 3: *Transition matrix*. Next period state variables, x_{t+1}, are independent on the realization of this period unobservable state variables, ε_t. The support of the observable state variables is finite and discrete. The transition across periods follows a first order Markov process.

Assumption 4: *Flow utility*. The flow utility of action 1 (continue to next period) is zero. Action 0 is a terminating action.

Assumption 5: *Discount factor*. The discount factor is known ($\beta \in (0, 1)$).

Nonparametric identification of the utility function is obtained employing a contraction mapping argument, given that the transiction matrix and the CCP are known. Therefore, $\Delta v(x_t) = v^1(x_t) - v^0(x_t)$ is known (obtained by inverting the CCP as shown in the main text - Section 4). Also, as a reminder, the *alternative specific value functions* are defined by

$$v^d(x_t) = \begin{cases} 0 + \beta \mathbb{E}^d\{v_{t+1}(x_{t+1})|x_t\} & \text{if } d = 1 \text{ (keep)} \\ u^0(x_t) & \text{if } d = 0 \text{ (sell)} \end{cases}$$

Step 1: Define the function $\hat{H}(r^0, r^1|x_t) = \mathbb{E}[\max_{d\in\{0,1\}}\{r^d + \varepsilon^d\}|x_t]$. Under the distributional assumption on the error term[7], $\hat{H}(\cdot|x_t)$ exists and has the additive property: $\hat{H}(r^0 + \kappa, r^1 + \kappa|x_t) = \hat{H}(r^0, r^1|x_t) + \kappa$ (see Rust (1994) and Magnac and Thesmar (2002)). This property is useful as it allows us to rewrite the *emax* function as the sum of a known object and an unknown function:

$$\hat{H}(v^0(x_t), v^1(x_t); x_t) = \hat{H}(v^0(x_t) - v^1(x_t), 0|x_t) + v^1(x_t) \equiv \hat{H}(\Delta v(x_t), 0|x_t) + v^1(x_t)$$

where $\hat{H}(\Delta v(x_t), 0|x_t)$ is identified because the difference in value function, $\Delta v(x_t) = v^0(x_t) - v^1(x_t)$, and the distribution of the error term, $\Lambda(\cdot)$, are known. To simplify the notation set $\hat{H}(x_{t+1}) = \hat{H}(\Delta v(x_t), 0|x_t)$.

Step 2: The alternative specific value function when the participant chooses to keep the asset, $v^1(x_t)$, is the unique solution of a functional equation. The following Lemma proves that $v^1(x_t)$ is a contraction.

Lemma: Denote by $\mathcal{X}$ the space of the observables and by $\mathcal{C}(\mathcal{X})$ the Banach space of all continuous, bounded functions $\omega : \mathcal{X} \to \mathbb{R}$. And define the operator $\Gamma : \mathcal{C}(\mathcal{X}) \to \mathcal{C}(\mathcal{X})$ by:

$$\Gamma\omega(x) = \beta\mathbb{E}^1\{\omega(x_{t+1})|x_t\}$$

Then, under the supremum norm, $||\omega|| = \sup_{x\in X} |\omega(x)|$, Γ is a contraction mapping with modulus β.

Proof: For any two functions $\omega, \hat{\omega} \in \mathcal{C}(\mathcal{X})$, we need to establish that $||\Gamma\omega - \Gamma\hat{\omega}|| \leq \mu||\omega - \hat{\omega}||$, for $\mu \in (0, 1)$. First, rewrite $\mathbb{E}^1\{\omega(x_{t+1})|x_t\} = \mathbb{E}[\max_{d\in\{0,1\}}\{\omega^d(x_{t+1}) + \varepsilon^d\}|d_t = 1, x_t] = \mathbb{E}\{H(x_{t+1}) +$

[7] The error term, $\varepsilon_t = (\varepsilon_t^0, \varepsilon_t^1)$, has support on $\mathbb{R}^2$ and finite expectation $\mathbb{E}[\varepsilon^d] < \infty$ for $d \in \{0, 1\}$.

$\omega^1(x_{t+1})|d = 1, x_t\}$ by using the derivation in the first step. Then proceed as follows:

$$\begin{aligned}
||\Gamma\omega - \Gamma\hat{\omega}|| &= \sup_{x_t \in \mathcal{X}} \Big| \beta \mathbb{E}\{H(x_{t+1}) + \omega^1(x_{t+1})|d_t = 1, x_t\} - \beta \mathbb{E}\{H(x_{t+1}) + \hat{\omega}^1(x_{t+1})|d_t = 1, x_t\} \Big| \\
&= \sup_{x_t \in \mathcal{X}} \beta \Big| \mathbb{E}\{H(x_{t+1}) + \omega^1(x_{t+1}) - H(x_{t+1}) - \hat{\omega}^1(x_{t+1})|d_t = 1, x_t\} \Big| \\
&\leq \beta \sup_{x_{t+1} \in \mathcal{X}} \Big| \omega^1(x_{t+1}) - \hat{\omega}^1(x_{t+1}) \Big| \\
&= \beta ||\omega - \hat{\omega}||
\end{aligned}$$

Therefore Γ is a contraction mapping with modulus β. The second line moves the arguments from the second expectation to the first. The third line removes the equal terms ($H(x_{t+1})$) and the conditional expectation ($\leq$ follows from this). The fourth line is from the definition of the supremum norm. ■
Therefore, $v^1(x_t)$ is identified.

Step 3: In the previous steps we identified nonparametrically $\Delta v(x_t) = v^0(x_t) - v^1(x_t)$ (directly from the data), and $v^1(x_t)$ (by the Contraction Mapping Theorem). Therefore, $v^0(x_t) = \Delta v(x_t) + v^1(x_t)$ and because $v^0(x_t)$ consists only of the flow utility (it corresponds to the terminating action), then $u^0(x_t) = \Delta v(x_t) + v^1(x_t)$. ■

G Supplementary Graphs

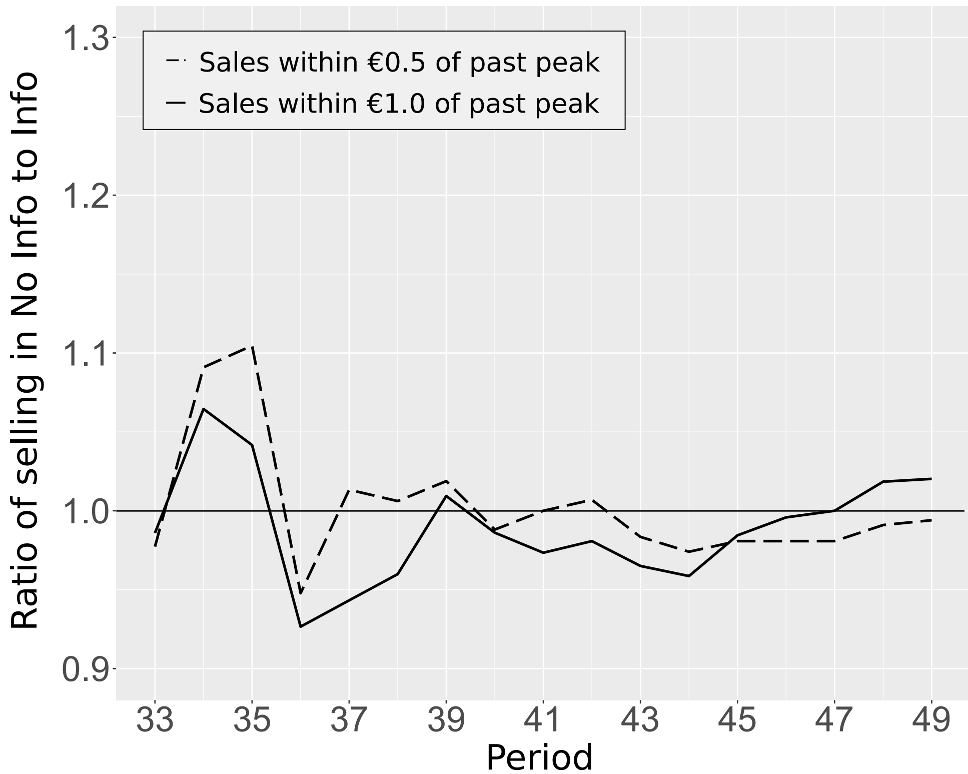

Figure 13: The ratio of the number of sales up to period t in the No Info to Info condition starting from period 33.

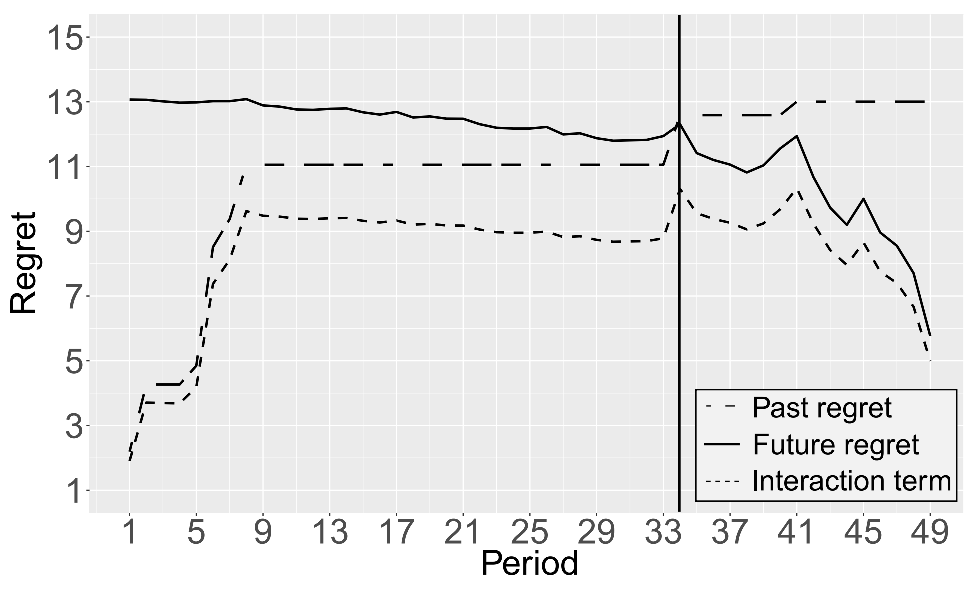

Figure 14: Example of the dynamic of past and future regret in two selected markets in the Info condition for the periods $t \in \{16, ..., 48\}$. The curves show the terms of the estimated regret function (i.e., past regret = $\hat{\omega}_I \, U(s_{p,t}; \rho)$, future regret = $\hat{\alpha}_I \, U(s_{f,t}; \rho)$ and interaction term = $|\hat{\lambda}_I| \, U(s_{f,t}; \rho) \, U(s_{p,t}; \rho)$). The solid vertical line shows the moment at which the participants switch the focus from future regret to past regret.

H 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	0.50	0.00	0.50	$\{0,1\}$	1 if the market condition is Info and 0 if it is No Info
period	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
entry price	5.30	5.18	1.50	$[2.30,8.56]$	Entry price in period 15
future expected price	5.00	5.00	0.08	$[3.48,7.02]$	Expected future price (over all remaining periods) conditional on the current price
past peak	7.58	7.52	0.59	$[5.53,8.56]$	Highest price in the past
future expected peak	7.64	7.79	0.45	$\{0\} \cup [3.48,8.23]$	Highest expected future peak conditional on the current price. 0 if the condition is No Info
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

Table 4: Variables used in the regression Tables 1 (Section 5), 5 and 6 (additional tables are reported in Appendix I). The statistics refers to all periods when a choice is made (periods 16 to 49).

I Additional Regressions

Pr[choice = keep]	I	II	III	IV
period	–0.104***	–0.103***	–0.105***	–0.100***
	(0.004)	(0.004)	(0.004)	(0.004)
price	0.116	0.159	0.291*	0.220
	(0.135)	(0.154)	(0.147)	(0.149)
price^2	–0.166***	–0.172***	–0.185***	–0.181***
	(0.012)	(0.014)	(0.014)	(0.014)
entry price	0.464***	0.438***	0.422***	0.424***
	(0.024)	(0.028)	(0.028)	(0.028)
future expected price	0.626*	0.386	0.417	0.429
	(0.274)	(0.308)	(0.304)	(0.277)
past peak			0.246***	0.371***
			(0.045)	(0.051)
future expected peak				0.269***
				(0.051)
past peak×info				–0.277***
				(0.051)
hl		–0.866*	–0.861*	–0.865*
		(0.403)	(0.403)	(0.403)
constant	5.229***	7.066***	4.897**	4.039**
	(1.321)	(1.481)	(1.507)	(1.376)
N	112,137	89,951	89,951	89,951

Table 5: 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.
***,** ,* denote statistical significance at the 0.1, 1, and 5 percent level.

	Early Markets		Late Markets	
Pr[choice = keep]	I	II	III	IV
period	–0.115***	–0.110***	–0.113***	–0.108***
	(0.004)	(0.004)	(0.005)	(0.005)
price	0.251	0.180	0.529**	0.451*
	(0.166)	(0.172)	(0.183)	(0.187)
price^2	–0.171***	–0.166***	–0.223***	–0.219***
	(0.015)	(0.015)	(0.016)	(0.016)
entry price	0.472***	0.473***	0.485***	0.486***
	(0.029)	(0.029)	(0.031)	(0.031)
future expected price	0.492	0.487⁻	0.817⁻	0.779*
	(0.306)	(0.283)	(0.428)	(0.394)
past peak	0.139**	0.249***	0.277***	0.401***
	(0.049)	(0.055)	(0.055)	(0.062)
future expected peak		0.249***		0.266***
		(0.062)		(0.070)
past peak×info		–0.250***		–0.278***
		(0.063)		(0.072)
constant	4.639**	3.943**	1.820	1.251
	(1.502)	(1.385)	(2.090)	(1.895)
N	60,373	60,373	51,764	51,764

Table 6: Random effects logit regression of the choice to keep the asset in early (1 to 24) and late (25 to 48) markets. 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). Errors are clustered by participant.
$^{***}, ^{**}, ^{*}, ^{-}$ denote statistical significance at the 0.1, 1, 5, and 10 percent level.

J Additional Estimations of the Structural Model

J.1 Different Regret Functions

This section reports estimates for five different models displaying different parameterization of the regret averse utility function. All estimations are consistents with the findings displayed in Section 6. The table show NLS estimates assuming the following discount rates: $\beta \in \{0.99, .98, .97\}$. In all models the objective function is (4.3) in Section 4, while the utility function is $u(y_t, s_{p,t}, s_{f,t}) = U(y_t; \rho) - R(s_{p,t}, s_{f,t}; \rho)$, where $U(y_t; \rho)$ is CRRA with risk aversion parameter ρ, and $R(\cdot, \cdot; \rho)$ is the regret function. While $U(\cdot; \rho)$ is constant across the five models, the regret function is not. The list below shows the regret functions employed for each model (numbered from 1 to 5):

1. $R = \omega\, U(s_{p,t}; \rho) + \alpha\, U(s_{f,t}; \rho)$
2. $R = \mathbb{1}_{\{\text{No Info}\}}\, \omega_{NI}\, U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} \left(\omega_I\, U(s_{p,t}; \rho) + \alpha_I\, U(s_{f,t}; \rho) \right)$
3. $R = \mathbb{1}_{\{\text{No Info}\}}\, \omega_{NI}\, U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} \left(\omega_I\, U(s_{p,t}; \rho) + \alpha_I\, U(s_{f,t}; \rho) + \lambda_I\, U(s_{p,t} \times s_{f,t}; \rho) \right)$
4. $R = \mathbb{1}_{\{\text{No Info}\}}\, \omega_{NI}\, U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} \left(\omega_I\, U(s_{p,t}; \rho) + \alpha_I\, U(s_{f,t}; \rho) + \lambda_I\, U(s_{p,t}; \rho) \times U(s_{f,t}; \rho) \right)$
5. $R = \mathbb{1}_{\{\text{No Info}\}}\, \omega_{NI}\, U(s_{p,t}; \rho) + \mathbb{1}_{\{\text{Info}\}} \left(\omega_I\, U(s_{p,t}; \rho) + \alpha_I\, U(s_{f,t}; \rho) + \lambda_I\, s_{p,t} \times s_{f,t} \right.$

The main specification is Model 4, because in this version the interaction term captures the cross-partial of the regret function across the past and future peaks, while controlling for risk attitudes (the estimates of Model 4 are reported in the main text in Table 2). Overall, the results are very similar across all tables, and corroborate our conclusions outlined in Section 7.

The dataset is discretized over 200 points in the interval [0.5, 9.3] according to the procedure laid out in Section D. The distance between adjacent points is €0.049. Section J.2 in the Appendix estimates models 1 to 5 over a support consisting of 300 points.

Parameter	$\beta = 0.99$	$\beta = 0.98$	$\beta = 0.97$
Model 1			
$\hat{\rho}$	-0.333^{***}	-0.334^{***}	-0.334^{***}
	(0.002)	(0.003)	(0.002)
$\hat{\omega}$	1.029^{***}	1.018^{***}	1.006^{***}
	(0.055)	(0.029)	(0.019)
$\hat{\alpha}$	0.213^{***}	0.168^{***}	0.135^{***}
	(0.025)	(0.021)	(0.018)
Model 2			
$\hat{\rho}$	-0.333^{***}	-0.334^{***}	-0.335^{***}
	(0.002)	(0.002)	(0.002)
$\hat{\omega}_{NI}$	1.106^{***}	1.002^{***}	1.001^{***}
	(0.098)	(0.034)	(0.022)
$\hat{\omega}_I$	1.106^{***}	1.029^{***}	0.977^{***}
	(0.098)	(0.064)	(0.053)
$\hat{\alpha}_I$	0.194^{***}	0.164^{***}	0.153^{***}
	(0.033)	(0.034)	(0.035)
N	111,613	111,613	111,613

Table 7: Estimation of models 1 and 2. Periods: 16 to 48. Standard errors are in parenthesis. ***,** ,* denote statistical significance at the 0.1, 1 and 5 percent level.

Parameter	$\beta = 0.99$	$\beta = 0.98$	$\beta = 0.97$
Model 3			
$\hat{\rho}$	-0.334^{***} (0.002)	-0.335^{***} (0.002)	-0.335^{***} (0.002)
$\hat{\omega}_{NI}$	0.988^{***} (0.068)	1.006^{***} (0.034)	1.002^{***} (0.022)
$\hat{\omega}_I$	1.268^{***} (0.102)	1.119^{***} (0.066)	1.062^{***} (0.056)
$\hat{\alpha}_I$	1.197^{***} (0.171)	0.856^{***} (0.140)	0.707^{***} (0.118)
$\hat{\lambda}_I$	-0.068^{***} (0.011)	-0.047^{***} (0.009)	-0.039^{***} (0.008)
Model 4			
$\hat{\rho}$	-0.334^{***} (0.002)	-0.335^{***} (0.002)	-0.335^{***} (0.002)
$\hat{\omega}_{NI}$	0.988^{***} (0.068)	1.006^{***} (0.034)	1.002^{***} (0.022)
$\hat{\omega}_I$	1.200^{***} (0.099)	1.072^{***} (0.064)	1.023^{***} (0.054)
$\hat{\alpha}_I$	1.129^{***} (0.160)	0.809^{***} (0.131)	0.668^{***} (0.110)
$\hat{\lambda}_I$	-0.090^{***} (0.015)	-0.063^{***} (0.012)	-0.052^{***} (0.011)
Model 5			
$\hat{\rho}$	-0.334^{***} (0.003)	-0.335^{***} (0.003)	-0.336^{***} (0.003)
$\hat{\omega}_{NI}$	0.986^{***} (0.068)	1.006^{***} (0.034)	1.003^{***} (0.022)
$\hat{\omega}_I$	1.271^{***} (0.120)	1.167^{***} (0.088)	1.132^{***} (0.077)
$\hat{\alpha}_I$	0.700^{**} (0.217)	0.583^{**} (0.188)	0.588^{***} (0.162)
$\hat{\lambda}_I$	-0.122^{*} (0.052)	-0.102^{*} (0.045)	-0.107^{**} (0.039)
N	111,613	111,613	111,613

Table 8: Estimation of models 3 to 5. Periods: 16 to 48. Standard errors are in parenthesis. ***,**,* denote statistical significance at the 0.1, 1 and 5 percent level.

J.2 Estimation on a Finer Support

This section provides robustness results showing that the discretization specified in Appendix D and based on 200 points does not affect the interpretation of the estimated coefficients. In particular, we report the previous analysis assuming that the current price takes value over a support consisting of 300 points. This considerably increases the difficulty of the minimization problem at hand as the size of the state space is now $300 \times 300 \times 300$. This is also reflected in much slower computing speed. The first five columns of Table 9 show the estimations for Models 1 - 5 as in Appendix J.1 while column six estimates the same specification as in Table 3 where future regret is allowed to affect not only the Info condition but also the No Info condition (called "Model 6"). The discount factor is $\beta = 0.99$.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
$\hat{\rho}$	-0.342^{***}	-0.341^{***}	-0.343^{***}	-0.343^{***}	-0.344^{***}	-0.347^{***}
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
$\hat{\omega}$	1.113^{***}					
	(0.059)					
$\hat{\alpha}$	0.245^{***}					
	(0.027)					
$\hat{\omega}_{NI}$		1.026^{***}	1.040^{***}	1.040^{***}	1.041^{***}	0.627^{***}
		(0.073)	(0.073)	(0.073)	(0.074)	(0.120)
$\hat{\omega}_I$		1.252^{***}	1.436^{***}	1.362^{***}	1.498^{***}	1.350^{***}
		(0.106)	(0.110)	(0.107)	(0.129)	(0.124)
$\hat{\alpha}_{NI}$						0.286
						(0.184)
$\hat{\alpha}_I$		0.210^{***}	1.329^{***}	1.255^{***}	0.960^{***}	1.286^{***}
		(0.035)	(0.180)	(0.168)	(0.230)	(0.195)
$\hat{\lambda}_{NI}$						-0.006
						(0.017)
$\hat{\lambda}_I$			-0.074^{***}	-0.100^{***}	-0.183^{**}	-0.101^{***}
			(0.012)	(0.016)	(0.056)	(0.018)
N	111,613	111,613	111,613	111,613	111,613	111,613

Table 9: Estimation results when the state space is discretized over 300 points. Periods: 16 to 48. Standard errors are in parenthesis.
***,** ,* denote statistical significance at the 0.1, 1 and 5 percent level.

J.3 Estimation with Heterogeneous Risk Preferences

When participants are endowed with heterogeneous risk preferences they use different cut-off rules which leads to different optimal selling times (prices). This could bias the results in Section 6. To investigate this issue we leverage on the overidentification of the problem at hand and estimate the same model as in Table 3 while assuming that ρ is a binomial random variable

$$\rho = \begin{cases} \rho^a & \text{with probability } \pi \\ \rho^b, & \text{with probability } 1 - \pi, \end{cases} \tag{J.1}$$

This is equivalent to assuming that agents belong to two types with different risk attitudes. In particular, we propose two different specifications. In Estimation 1 (first two columns on Table 10) both ρ^a and ρ^b are allowed to vary freely on the real line. The estimation results mirror those in Table 3, with $\hat{\rho}^a = \hat{\rho}^b \approxeq -0.3$: this implies that there is only one type of risk preferences.

The main problem in our setting arises when some participants have concave utility functions while others have convex ones. In Estimation 2 we constraint ρ^a to be negative and ρ^b to be non-negative. In Table 10 ρ^b is estimated to be at the boundary of its domain, and both risk aversion parameters are not significantly different from zero. The estimates of the other parameters support our previous finding.

	Estimation 1		Estimation 2	
Parameter	$\beta = 0.99$	$\beta = 0.97$	$\beta = 0.99$	$\beta = 0.97$
$\hat{\rho}^a$	-0.322	-0.321	-0.359	-0.336
	(...)	(...)	(1.359)	(1.703)
$\hat{\rho}^b$	-0.323	-0.321	0.000	0.000
	(...)	(...)	(...)	(...)
$\hat{\pi}$	0.739	0.633	0.867	0.940
	(...)	(...)	(7.912)	(...)
$\hat{\omega}_{NI}$	0.640***	0.857***	0.637***	0.855***
	(0.120)	(0.065)	(0.119)	(0.065)
$\hat{\omega}_I$	1.242***	1.100***	1.247***	1.102***
	(0.125)	(0.069)	(0.126)	(0.078)
$\hat{\alpha}_{NI}$	-0.153	-0.177	-0.117	-0.040
	(0.190)	(0.131)	(0.189)	(0.131)
$\hat{\alpha}_I$	0.848**	0.412*	0.945***	0.724***
	(0.271)	(0.190)	(0.281)	(0.210)
$\hat{\lambda}_{NI}$	-0.001	0.008	-0.004	-0.006
	(0.018)	(0.013)	(0.017)	(0.012)
$\hat{\lambda}_I$	-0.085**	-0.045*	-0.092*	-0.074*
	(0.027)	(0.019)	(0.036)	(0.031)
N	111,613	111,613	111,613	111,613

Table 10: The estimation of (4.3) with the regret term (6.5) in periods 16 to 48 for different values of the discount factor β. Standard errors are in parenthesis. The CCP is computed using (6.2) for both conditions.
***,**,* denote statistical significance at the 0.1, 1, and 5 percent level.
(...) denotes standard error above 10.

K Instructions (English)

K.1 Market Task

K.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**.

K.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.

K.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).

K.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

K.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).

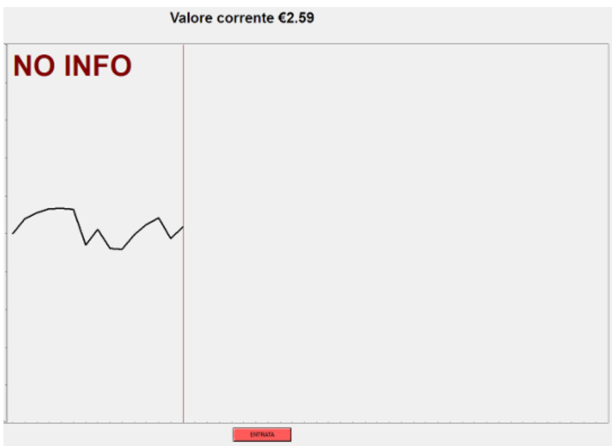

K.1.6 Exiting the Market

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.

K.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 luscita” (Information after exit) or “NO INFO aluscita” (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.

K.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.

K.1.9 Payment

You 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.

K.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.**

L Instructions (Italian)

L.1 Market Task

L.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**.

L.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.

L.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 pu 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**).

L.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

L.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.”**

L.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.

L.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 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.

L.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.”**

L.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.

L.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.**

L.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 presenata 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.

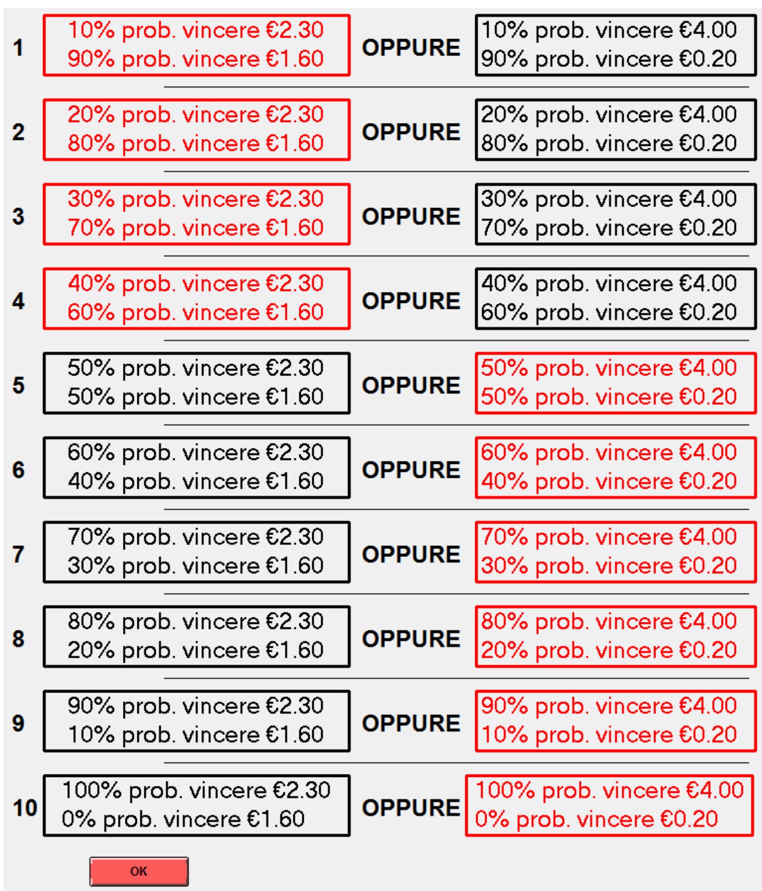

M Supplementary Regressions

M.1 Analysis of early sales in the info condition

Pr[choice = keep]	I early: 20	II early: 23	III early: 25	IV early: 28	V early: 30
period	–0.104***	–0.113***	–0.126***	–0.131***	–0.142***
	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)
price	0.239*	0.188	0.111	0.155	0.213
	(0.132)	(0.131)	(0.133)	(0.133)	(0.131)
price^2	–0.178***	–0.174***	–0.169***	–0.172***	–0.177***
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
entry price	0.448***	0.454***	0.457***	0.443***	0.441***
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
future expected price	0.614**	0.739***	0.892***	0.892***	0.911***
	(0.270)	(0.259)	(0.245)	(0.246)	(0.245)
past peak	0.217***	0.223***	0.219***	0.211***	0.221***
	(0.038)	(0.038)	(0.038)	(0.038)	(0.038)
info	–0.077*	–0.096**	–0.125**	–0.155***	–0.092
	(0.040)	(0.047)	(0.051)	(0.055)	(0.061)
early	–0.028	–0.251***	–0.542***	–0.635***	–0.832***
	(0.064)	(0.065)	(0.071)	(0.067)	(0.070)
info×early	0.194***	0.173***	0.207***	0.227***	0.109
	(0.061)	(0.055)	(0.059)	(0.064)	(0.067)
constant	3.465***	3.283***	3.303***	3.516***	3.697***
	(1.328)	(1.254)	(1.163)	(1.153)	(1.127)
info's marginal effect	0.116**	0.083*	0.082*	0.072*	0.017
	(0.059)	(0.045)	(0.043)	(0.044)	(0.041)
N	112,137	112,137	112,137	112,137	112,137

Table 11: The regression table mirrors numerically the intuition in Panel B of Figure 2 as participants in info sell less often early on because of the possibility of future regret. In fact, the last panel shows that the marginal effect of info decreases as the variable early includes more periods. 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.
Random effects logit regression of the choice to keep the asset. 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). Errors are clustered by participant. The descriptions of all the variables (other than entry) can be found in Appendix H.
***,** ,* denote statistical significance at the 1, 5 and 10 percent level.

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. and VOSTROKNUTOV, A. (2016). The impact of the degree of responsibility and mutual decision making on choices under risk, 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. *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.

MURPHY, A. (2015). A dynamic model of housing supply. *Working Paper SSRN 2200459*.

POTUSCHAK, H. and MULLER, W. G. (2009). More on the distribution of the sum of uniform random variables. *Statistical Papers*, **50** (1), 177–183.

RUST, J. (1994). *Handbook of Econometrics*, Amsterdam. North Holland, vol. 4, chap. Structural estimation of Markov decision processes, pp. 3081–3143.

TAUCHEN, G. (1986). Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, **20**, 177–181.