

Ignorance is Bliss: A Game of Regret*

Claudia Cerrone[†]

Francesco Feri[‡]

Philip R. Neary[§]

September 24, 2021

Abstract

The outcome of a foregone alternative is not always learnt. We incorporate this observation into a model of regret, supposing that the ex-post information available depends on choice. We show that a more informative ex-post environment is never desirable for a regret averse individual. We then suppose that there are multiple regret averse individuals where the ex-post information available depends on own choice and the choices of others. This implies that what appears to be a series of isolated decision problems is in fact a behavioural game with multiple equilibria. We experimentally test the model and find support for our theory.

JEL codes: C72; C92; D81; D91.

Keywords: regret aversion; coordination games; information.

*Thanks to Zvonimir Bašić, Phil Brookins, Christoph Engel, Andreas Grünewald, Evan Piermont, Antonio Rosato, and audiences at University of Tilburg (TIBER 2019), UCSD (Southwest Economic Theory conference 2020), UCSB, King's College London and University of Bath. Funding from the Max Planck Society, Middlesex University and Royal Holloway University is gratefully acknowledged.

[†]*Email:* c.cerrone@mdx.ac.uk *Address:* Middlesex University London, The Burroughs, Hendon, London, NW4 4BT. [corresponding author]

[‡]*Email:* francesco.feri@rhul.ac.uk; *Address:* Department of Economics, Royal Holloway, University of London, Egham, Surrey, TW20 0EX.

[§]*Email:* philip.neary@rhul.ac.uk; *Address:* Department of Economics, Royal Holloway, University of London, Egham, Surrey, TW20 0EX.

For of all sad words of tongue or pen,
The saddest are these: 'It might have
been!'

John Greenleaf Whittier

1 Introduction

Ann decides to go to her local Italian restaurant for dinner. Ann whittles down her choice to either the spaghetti or the risotto. The spaghetti is always the same and always pretty good - a safe choice. The risotto is risky, sometimes it is excellent but on occasion it has been bad. Ann suffers from regret. If she chooses the risotto and it turns out to be bad then she will regret her choice; likewise, if she chooses the spaghetti and the risotto turns out to be excellent. If Ann is regret averse, then she will account for the possibility of these unpleasant regretful experiences when making her decision.

Information plays an integral role in the above story. Ann's regret is evaluated as the utility loss experienced from comparing a choice made - that turned out to be suboptimal - to a foregone alternative. But implicit in this construction is that such a comparison can always be made. While this may be true for financial assets listed on an exchange, there are many situations in life where information about unchosen options is not automatically available, as in the case of our restaurant example as it is uncommon to order two main courses. In this paper we allow for this possibility, by allowing for the ex-post information available to a decision maker to depend on choice.

Now suppose the environment is modified slightly in that if Ann orders the spaghetti then she will never learn the quality of the risotto, but since the spaghetti is always the same its quality is learned whether it is chosen or not. In this environment there is an informational asymmetry. Does this matter? The answer is: yes, it can matter because differing levels of ex-post information matter for behavioural individuals like Ann who factor in ex-post information into their decision-making. In this new environment, the spaghetti dish is now more relatively desirable than before because choosing it completely insures against regret, while the total benefit associated with choosing the risotto remains unchanged. Depending on the parameters, it is easy to see that ex-post information structures varying with choice can lead to preference reversals.

People may learn from others. Imagine that Ann has a friend called Barry who joins her for dinner. If Barry orders the risotto, then ordering the spaghetti is no longer

a safe haven from regret for Ann as it was when she was dining alone. Previously by ordering the spaghetti, Ann was completely insured against regret. But now, if Ann orders spaghetti and learns via Barry’s that the risotto was excellent, then she will experience regret. In this case, Barry’s choice has imposed an externality on Ann, because his ordering the risotto Barry will impact Ann’s ex-post information, which in turn impacts the psychological payoff (i.e., the payoff from the choice made plus the potential psychological loss due to regret) she associates with a given choice. So, what on the surface appeared to be two independent decision problems made in isolation is in fact a behavioural game since each person’s choice can reveal information to others.

In this paper we present a model that captures the Ann and Barry story above.¹ While the model can handle environments far richer than this example, we believe that the simple story captures all the salient features. That is, we start out with the (obvious) observation that ex-post information matters for regret. We then provide some formal machinery that allows the analyst to model how the ex-post information available to a regret averse individual can vary with choice and hence affect optimal choice.² Effectively, this requires extending the domain of preferences from ‘objects of choice’ to ‘objects of choice and their associated information environment’. Our first result, Theorem 1, shows that when there is full information ex-post information, an individual with standard preferences and an individual who is regret averse with a linear regret term are observationally equivalent in their choice behaviour.

We then observe how Theorem 1 breaks down once ex-post information depends on choice. In our opinion it seems important that experimenters are aware of this fact. A full classification of precisely how and when ex-post information environments impact optimal choice is not possible: not all information sets can be ordered. But our second result, Theorem 2, shows how information sets can be partially ordered. In particular, we show that ‘more’ ex-post information (formally defined later) is never desirable for a regret averse decision maker. Ignorance is bliss. In a sense this is obvious: a regret averse individual ‘zeroes-in’ on the best-performing lottery that is learned about in each

¹The example is based on [Ariely and Levav \(2000\)](#). In that paper, the focus was on the differences in patrons’ orders in two different settings: one where patrons order simultaneously and another where they order sequentially. [Ariely and Levav \(2000\)](#) noted disparities in order choice (i.e., the propensity to “coordinate”) depending on the protocol. The model in this paper can explain this finding.

²In fact we allow for the possibility that the ex-post information structure is not fixed. That is, there may be multiple ex-post informational environments associated with a given choice and there is a given probability distribution over the different ex-post information structures.

state, creating a state-dependent reference point that can only be matched but never exceeded. So when the number of outcomes that are learned about increases, regret can only go up.

After this, we allow for the possibility that the ex-post information available depends not only on one’s own choice but also on the choices of others. Precisely, we consider an environment wherein a large number of regret averse individuals choose from a common choice set. We show that such an environment is not a series of independent decision problems to be analysed in isolation, but is in fact a rich multi-player behavioural game, that we term the regret game. When the behaviour of others can impact one’s ex-post information, for certain parameter specifications the regret game is a game of coordination that admits multiple equilibria; this despite the fact that the same individuals all have a strictly dominant choice when faced with the same decision problem in isolation. Theorem 3 is a precise statement of this.

The term ‘asymmetric information’ is typically understood to mean ‘ex-ante asymmetric information’. The idea is that in advance of some economic interaction, say a potential exchange, one individual is relatively more informed. The focus is on how the outcomes in such an incomplete information environment may differ from that wherein information is complete. The regret game has, in a sense, exactly the opposite core feature of ‘ex-post asymmetric information’. At the onset everyone possesses the same information; but after actions have been taken individuals may possess differing levels of information. A particular type of market failure can emerge: all individuals take a decision that is collectively suboptimal so as to avoid being the one with more, not less, information ex-post.

We then move to testing our theory experimentally. In a big-picture sense, we begin by eliciting certainty equivalents when ex-post information is total and compare it with the elicited certainty equivalents when ex-post information depends on choice. We find strong support for Theorem 2 in that many subjects make different choices when the only difference is the ex-post information supplied. (Theorem 1 ensures that our mechanism is incentive compatible.) The results from these initial decisions also allow us to classify participants into those who are regret averse and those who are not,³

³As further discussed in Section 3.1, the idea that the expectation of feedback affects anticipated regret and thus behaviour has been studied in the psychology literature (Zeelenberg et al., 1996). Yet ours is the first paper that identifies regret averse individuals by manipulating feedback and using an incentive-compatible mechanism.

and to calibrate the parameters used in the second part of the experiment wherein our participants play a version of the regret game.

More specifically, we run two experiments both using a within-subject design. The first experiment consists of two parts. In the first part, we elicit participants' preferences over a sure amount of money and a risky lottery under two different information conditions – one where participants learn the risky lottery's outcome even if they do not choose it, and one where they do not learn the risky lottery's outcome unless they choose it. In the second part of the experiment, participants are matched in pairs and play the regret game. They must choose between a sure amount of money and a lottery. If they do not choose the lottery, they will learn its outcome only if their partner chose the lottery.

We find that, when the regret game has a unique equilibrium in dominant strategies, the vast majority of participants choose it whether they are regret averse or not. When the regret game is a game of coordination, regret averse participants try to coordinate with their partner. This supports our model's predictions. We observe a positive and highly significant impact of beliefs on choice for regret averse participants. When we focus on the last iteration of the game, the effect is even stronger. We also observe a positive impact of beliefs on choice for non regret averse participants, but the effect is smaller and only marginally significant.

These results indicate that regret aversion can drive coordination. However, in our setting there may have been two additional drivers of coordination: preferences for conformism ([Charness et al., 2017, 2019](#)) and inequity averse preferences ([Fehr and Schmidt, 1999](#)). To eliminate these potential confounds we ran a second experiment. The second experiment consists of two parts. In the first part, participants choose between a sure amount of money and a risky lottery, with subjects revealing to us their certainty equivalent for the risky lottery. Subjects are later asked whether they want to find out how much they would have earned had they chosen the risky lottery. If they choose not to find out, they forgo a small amount of money. This procedure allows us to classify participants into those who are regret averse and those who are not. In the second part of the experiment, participants are matched in pairs and play a variant of the regret game: they must choose whether they want to find out the risky lottery's outcome or not, but they can avoid finding out only if their partner also decides not to find out.

Again we find that regret averse participants try to coordinate with their partner. The probability that they choose what they think their partner chose is significantly higher than the probability that they choose the alternative. All regret averse participants who believe that their partner chose to avoid information, choose to avoid information. In contrast, non regret averse participants play the dominant strategy, and some of them do so under the belief that their partner chose the alternative option. Thus, once the aforementioned confounds are removed by design, we still find strong support for our model’s predictions.

Regret aversion dates back to the classic works of [Bell \(1982\)](#) and [Loomes and Sugden \(1982\)](#). However, these and subsequent studies of regret aversion assume that ex-post information is total. In order for any behavioural bias to “have bite” after a decision has been made, information about foregone alternatives is required. We show how to model behavioural agents with biases that “kick in” after decisions have been taken, and further to document how choice can vary with uncertainty over what environment will be faced ex-post (in addition to the standard uncertainty associated with the outcomes of each choice). While we have focused on regret, we are hopeful the machinery will be useful to address other biases too.

While a number of papers explore how regret impacts choice in individual decision problems, the impact of regret averse preferences in strategic environments is extremely limited. A notable exception is [Filiz-Ozbay and Ozbay \(2007\)](#) who incorporate anticipated regret into the preferences of agents partaking in a sealed-bid auction.⁴ However, unlike in our paper, individuals in their environment either *always* find out the outcome of the foregone alternative or never find out. That is, the informational environment is set exogenously and does not depend endogenously on the behaviour of individuals. In our set up, anticipated regret can serve as a coordination device.

In models of *social learning* ([Bikhchandani et al., 1992](#); [Banerjee, 1992](#)), the behaviour of others generates information but is payoff irrelevant. Mapping this to the Ann and Barry story, individuals choose sequentially basing their choice on a weighted combination of prior belief and the choices of those who went before. In a model of *social interaction* (a game), someone else’s choice is payoff relevant, but does not generate

⁴There is a large literature that studies the connection between dynamics based on “regret minimisation” and convergence to Nash equilibria and correlated equilibria ([Aumann, 1974](#)). See for example, [Foster and Vohra \(1997\)](#), [Hart and Mas-Colell \(2000\)](#), and the textbook [Cesa-Bianchi and Lugosi \(2006\)](#).

information. The regret game is a behavioural game that, in a sense, lies somewhere between the two settings above: someone else’s choice is payoff relevant, but only because it can affect ex-post information and this can be potentially harmful to a behavioural individual.

The two papers closest to ours are [Bénabou \(2013\)](#) and [Cooper and Rege \(2011\)](#). Using preferences for late-resolution of uncertainty from [Kreps and Porteus \(1978\)](#), [Bénabou \(2013\)](#) presents a dynamic model that addresses the harmful issue of “group-think”. In the model each individual’s payoff depends on the effort level of everyone (including his own) and the realisation of a random variable. The key feature is the inclusion of *anticipatory utility* experienced from thinking about one’s future prospects. The more positive the forecast, the better for the individual. This allows for multiple equilibria including, for example, one in which everyone in the population collectively ignores a negative public signal about the random variable. Such delusions can persist because individual j ’s informational decision and effort choice can affect the risks of individual i ’s psychological payoff, leading i to make a different informational decision than he otherwise would. In our model by contrast, j ’s choice of lottery can impact the informational environment that i faces, changing i ’s psychological payoffs and leading i to make a different risk-taking decision than he would in isolation.

[Cooper and Rege \(2011\)](#) present a model of “social regret” wherein the regret from a choice that turned out suboptimal is dampened if others chose the same. An individual tells himself that his decision could not have been too wrong if many others acted the same. Misery loves company. This model’s key feature is the belief that an individual assigns to others choosing an alternative. The more likely another individual is to choose an alternative, the greater the expected regret from not choosing that alternative. Results from laboratory experiments confirm their hypothesis that social regret is a powerful force. In our model, an individual’s choice imposes an externality on others by changing the information environment that they may face. In [Cooper and Rege \(2011\)](#), the effect of social regret is magnified by the number of others who make a different choice, but social regret is less intense when others make the same choice.⁵

The remainder of the paper is structured as follows. Section 2 presents the model. Section 3 describes the designs and results of our two experiments. Section 4 concludes.

⁵A related literature is that on envy or “keeping up with the Joneses”, according to which people care not only about the absolute value of their own consumption, but also about the average (or per capita) level of consumption in the economy ([Gali, 1994](#)).

2 The Model

Ex-post information feedback is an integral part of the regret story. If the outcome of a foregone alternative is never learned, then how can it be regretted? In subsection 2.1, we build on this insight to formally explore the mechanics of regret aversion. In effect, a regret averse individual ‘zeroes-in’ on the best-performing lottery in each state, creating a state-dependent reference point that can only be matched but never exceeded. Formally, identifying each state of nature by its best-performing lottery generates a *regret-relevant information partition* of the state space. However, implicit in this construction is that alternatives, in particular the best-performing lottery, in each state will be learned about even if unchosen. We consider how preference reversals may occur when the regret-relevant partition that is faced depends on choice. Using our example from Section 1, we imagine that Ann only learns the outcome of the risotto if she orders it (since the spaghetti dish is a sure thing, it is always known).⁶ This asymmetry in ex-post information, and hence regret, increases the relative benefit of choosing spaghetti as it is a safe haven from regret.

In Subsection 2.2, we extend the environment to one in which the information that will be learned ex-post is unknown in advance. That is, in addition to the ‘standard’ uncertainty captured by risky lotteries, there is an additional layer of uncertainty due to not knowing what regret-relevant information partition will be faced ex-post. Regret may or may not be experienced. Again using our example, we suppose that there is a possibility that Ann will learn of the risotto’s quality even if she orders the spaghetti (perhaps a she will observe someone at a different table order it). Ordering the spaghetti no longer provides full insurance from regret. In Subsection 2.3 we extend to an environment where there are many regret averse individuals, and the regret-relevant information partition that is faced by an individual depends probabilistically upon her own choice *and* the choices of others. We suppose that Ann has a friend, Barry, who joins her for dinner. If Barry orders the risotto, then Barry’s choice means that Ann will learn the risotto’s quality. In this case, Barry’s behaviour imposes a negative externality on Ann’s ex-post psychological payoff. This is a game in which Barry does not alter Ann’s payoff directly, but rather he impacts the ex-post informational environment that Ann will face.

⁶We emphasise that risk free lotteries therefore play a very important role for regret averse individuals.

2.1 The Decision Environment

Let Ω denote a finite state space with typical element given by ω . Uncertainty is captured by a probability measure \mathbb{P} defined on 2^Ω , where $\mathbb{P}[\{\omega\}] > 0$ for all $\omega \in \Omega$.⁷ There is a choice set, \mathcal{L} , containing n risky lotteries, labelled $\ell_1, \ell_2, \dots, \ell_n$, and a safe (risk-free) lottery, ℓ_S . The outcome of lottery ℓ in state $\omega \in \Omega$ is denoted by $\ell(\omega)$. For simplicity we assume that each risky lottery, ℓ_i , $i = 1, \dots, n$, is an independent Bernoulli random variable with outcome $\underline{\ell}_i$ occurring with probability $1 - p_i \in (0, 1)$ and outcome $\bar{\ell}_i$ occurring with probability p_i . We further assume that there are no payoff ties and that outcomes are structured such that $\max_i \underline{\ell}_i < \ell_S < \min_i \bar{\ell}_i$. Therefore there are exactly 2^n states in Ω . Note that the risk free lottery, ℓ_S , is the lottery with highest return in only one state; in all other states, at least one of the risky lotteries will outperform it.⁸

Let $u(\cdot)$ be a real-valued *choiceless utility* function defined on $\mathcal{L} \times \Omega$ that satisfies the usual conditions.⁹ Let $R(\ell; \omega)$ capture the experienced regret in state ω when lottery ℓ was chosen; formally it is a function of the difference between ‘choiceless utility from what turned out to be the best possible decision’ and ‘choiceless utility from the decision made’.¹⁰ The total utility experienced by a decision maker in state $\omega \in \Omega$, when lottery $\ell \in \mathcal{L}$ is chosen, is denoted by u^T and is defined as

$$\begin{aligned} u^T(\ell(\omega)) &= u(\ell(\omega)) - R(\ell; \omega) \\ &= u(\ell(\omega)) - \kappa \left(\max_{\ell' \in \mathcal{L}} u(\ell'(\omega)) - u(\ell(\omega)) \right) \end{aligned} \tag{1}$$

where parameter κ (≥ 0) is the coefficient of regret aversion that is assumed to be state-independent.

A decision maker compares the expected total utility of all the lotteries. Letting \mathbb{E} denote the expectation operator with respect to \mathbb{P} , we can state the decision maker’s

⁷Since Ω is finite there are no technical headaches.

⁸Since the n risky lotteries have uncorrelated returns, then the probability that the risk-free lottery is the best-performing lottery equals $\prod_{i=1}^n (1 - p_i)$. This tends to zero as n gets large.

⁹The term “choiceless utility” was introduced by [Loomes and Sugden \(1982\)](#). It is so-called as it is the utility experienced if the decision maker is simply *assigned* lottery ℓ and the resulting state is ω

¹⁰This is different to the set up of [Loomes and Sugden \(1982\)](#), who view regret as stemming from pairwise comparisons. Our formulation is based on [Sarver \(2008\)](#), where the comparison lottery is that which performed best in the realised state.

optimisation problem as,

$$\max_{\ell \in \mathcal{L}} \mathbb{E} \left[u^T(\ell(\omega)) \right] \quad (2)$$

We now state our first result.

Theorem 1. *Suppose the outcome of all lotteries will always be learned ex-post. Then, for every pair of lotteries ℓ_i and ℓ_j in \mathcal{L} ,*

$$\mathbb{E} \left[(u(\ell_i(\omega))) \right] \geq \mathbb{E} \left[(u(\ell_j(\omega))) \right] \iff \mathbb{E} \left[(u^T(\ell_i(\omega))) \right] \geq \mathbb{E} \left[(u^T(\ell_j(\omega))) \right]$$

The proof is in the Appendix. Theorem 1 states that two decision makers, one with standard preferences ($\kappa = 0$) and the other regret averse ($\kappa > 0$), are indistinguishable in their choice behaviour. While the proof is almost immediate, it is important for our experimental purposes as it ensures that the incentive compatible mechanism that we use in our experiments, the BDM (Becker et al., 1964), remains incentive compatible under regret averse preferences. We note that while Theorem 1 is stated for Bernoulli lotteries, it is easily extended to an environment with more general lotteries provided everything remains finite.

While interesting, Theorem 1 does require that the outcome of every lottery is learned ex-post. However, while a lottery's outcome will always be learned when it is chosen, the same may not be true for unchosen lotteries. In such environments, what really matters to a regret averse individual is not the best performing lottery in each state, but the best performing lottery in each state that is learned about. The standard regret framework is insufficient for these purposes and needs to be extended to allow for the possibility that the information available ex-post depends on choice. We develop such an extension now.

Begin by relabelling the lotteries in such a way that $\bar{\ell}_1 > \bar{\ell}_2 > \bar{\ell}_3 > \dots > \bar{\ell}_{n-1} > \bar{\ell}_n > \ell_S$. Now, labelling lottery S by $n + 1$, for each $i = 1, \dots, n + 1$, define the event $F(i)$ as follows:

$$F(i) := \left\{ \omega \in \Omega \mid \ell_i(\omega) = \max_{\ell_j \in \mathcal{L}} \ell_j(\omega) \right\} \quad (3)$$

In words, event $F(i)$ is the set of states upon which lottery ℓ_i is the best-performing. From our assumptions on the structure of lottery outcomes each $F(i)$ is nonempty and the collection $\{F(j)\}_{j=1}^{n+1}$ forms a partition of the state space Ω . In words, a regret-averse individual creates a reference point for every state, where the state-dependent

reference point is the outcome of the best-performing lottery. Such a reference point can only be matched but never exceeded.

The partition given in (3) assumes that ex-post information will be total. But what happens when the ex-post information the decision maker receives depends on choice? That is, consider the event $F(j)$ but suppose that the environment is changed such that the outcome of ℓ_j in state $\omega \in F(j)$ is not learned when lottery ℓ_k is chosen. Clearly the state dependent reference point $\ell_j(\omega)$ is no longer valid. We assume that it is replaced with the payoff to the best-performing lottery that is learned about in state ω .

When lottery ℓ_k is chosen, we let $\mathcal{O}_k \subseteq \mathcal{L}$ denote the set of lotteries whose outcomes are observed.¹¹ Clearly \mathcal{O}_k is non empty for all k as the outcome of the chosen lottery is always learned. To incorporate ex-post information that depends on choice, we amend (3) above, and define $F_k(j)$ as follows.

$$F_k(j) := \left\{ \omega \in \Omega \mid \ell_j(\omega) = \max_{\ell_h \in \mathcal{O}_k} \ell_h(\omega) \right\} \quad (4)$$

That is, $F_k(j)$ is the set of states where lottery j is the best performing lottery that is learned about conditional on lottery k being chosen. Unlike the events defined in (3), it is possible for $F_k(j)$ to be empty. This is despite the fact that every lottery is the best-performing in at least one state.

So, for every lottery k the decision maker associates a partition of Ω given by $\pi_k = \{F_k(j)\}_{\ell_j \in \mathcal{L}}$. We call this the *regret relevant information partition* associated with lottery k . Note that when the outcomes of all lotteries can be learned ex-post, $F_k(j) = F(j)$, for every lottery k .¹²

We define an *ex-post information environment*, Π , as the collection of regret relevant information partitions, one for each lottery. That is, $\Pi = \{\pi_k\}_{k=1}^{n+1}$. We now show how to define a partial order on the collection of all ex-post information environments, that allows us to rank them in terms of greater ex-post ‘informativeness’.

Definition 1. We say that ex-post information environment $\Pi' = \{\pi'_k\}_{k=1}^{n+1}$ is *more informative* than ex-post information environment $\Pi'' = \{\pi''_k\}_{k=1}^{n+1}$ if for every $k =$

¹¹Note that this assumes that, conditional on choosing lottery k , the set of lotteries whose outcomes are learned does not vary with the realised state. While this generalisation is easily incorporated, the simpler version is sufficient for our purposes.

¹²In this case the collection $\{F(i)\}_{i=1}^{n+1}$ forms the *fundamental regret relevant information partition* that we denote by π_0 .

$1, \dots, n + 1$, regret relevant information partition π'_k is as fine as regret relevant information partition π''_k , and π'_j is strictly finer than π''_j for at least one lottery j .¹³

We now show that a more informative ex-post information environment is not desirable for a regret averse individual.

Theorem 2. *Consider two ex-post information environments Π' and Π'' such that Π' is more informative than Π'' . Then a regret averse individual prefers ex-post information environments Π'' to Π' , in the sense that the expected total utility associated with a given choice of lottery is never higher in Π' .*

Theorem 2 is very intuitive. A regret averse individual constructs a reference point for every (lottery, state) pair, that is given by the best-performing lottery that is learned about. A higher reference point is bad as the difference between the reference point and the chosen lottery's outcome can only increase, thereby increasing the regret that is experienced. By the definition of more informativeness, the reference point is never decreasing and strictly increases for at least one (lottery, state) pair. Since the choiceless utility experienced does not depend on the ex-post information environment and a more informative environment brings with it more regret, then such an environment is not attractive to a regret averse individual.

We now show how to encode the example of Ann deciding between restaurants from Section 1 into our framework.

Example 1. Regret averse Ann must decide between two pizza-serving restaurants, R and S . There are two states of the world, ω_1 and ω_2 , with the probability of state ω_1 given by $p \in (0, 1)$. Restaurant R is risky and its pizza brings payoff \underline{l}_R in state ω_1 and payoff \bar{l}_R in state ω_2 . Restaurant S is the safe choice, with its pizza bringing payoff l_S in both states. Payoffs are structured such that $\underline{l}_R < l_S < \bar{l}_R$.

There are two information environments Π' and Π'' . Under Π' , the outcome of both lotteries is always learned ex post, while under Π'' , the outcome of lottery R is learned only if chosen. Therefore, we have

$$\begin{aligned} \text{Under } \Pi' : \quad & \mathcal{O}_{\underline{l}_R} = \{\underline{l}_R, l_S\}, \mathcal{O}_{l_S} = \{\underline{l}_R, l_S\} \text{ and } \pi'_{\underline{l}_R} = \{\{\omega_1\}, \{\omega_2\}\} = \pi'_{l_S} \\ \text{Under } \Pi'' : \quad & \mathcal{O}_{\underline{l}_R} = \{\underline{l}_R, l_S\}, \mathcal{O}_{l_S} = \{l_S\} \text{ and } \pi''_{\underline{l}_R} = \{\{\omega_1\}, \{\omega_2\}\} \neq \pi''_{l_S} = \{\{\omega_1, \omega_2\}\} \end{aligned}$$

¹³Given a set S and two partitions of S , ρ_1 and ρ_2 , we say that ρ_1 is finer than ρ_2 if every element α of ρ_1 is a subset of some element of ρ_2 .

Let us now consider the expected total utility associated with each choice of lottery. That is, we evaluate $u^T(\ell, \Pi)$, where we note the explicit dependence on the information environment Π . We have

$$\begin{aligned} u^T(\ell_R, \Pi') &= p\left(u(\underline{\ell}_R) - \kappa(u(\underline{\ell}_R) - u(\ell_S))\right) && +(1-p)(u(\bar{\ell}_R)) \\ u^T(\ell_R, \Pi') &= pu(\ell_S) && +(1-p)\left(u(\ell_S) - \kappa(u(\ell_S) - u(\bar{\ell}_R))\right) \end{aligned}$$

while

$$\begin{aligned} u^T(\ell_R, \Pi'') &= p\left(u(\underline{\ell}_R) - \kappa(u(\underline{\ell}_R) - u(\ell_S))\right) && +(1-p)(u(\bar{\ell}_R)) \\ u^T(\ell_R, \Pi'') &= pu(\ell_S) && +(1-p)u(\ell_S) \end{aligned}$$

where we note that under information environment Π'' , by choosing ℓ_S Ann is completely insured against regret. Note that S has become relatively more attractive to Ann in Π'' .

We conclude this section by reiterating the main observation. Models of regret aversion assume that the decision maker identifies each state by its best-performing lottery. But this is not always possible. When this is the case the regret averse decision maker identifies each state by the best-performing lottery that is learned about. Finally, what lottery outcomes are learned and what are not may be choice dependent. It then immediately follows that varying the information environment can impact (optimal) choice. In the next subsection we will allow for the possibility that multiple regret-relevant information partitions can be associated with each lottery choice. That is, there will be further uncertainty about what ex-post information environment will be faced.

2.2 Uncertainty Over the Ex-post Information Structure

In this section we allow for the possibility that the outcomes of unchosen lotteries may be learned about or may not be learned about. That is, there may be more than one regret relevant partition associated with a lottery choice. To illustrate the effect this can have on optimal choice we limit attention to the environment where the choice set \mathcal{L} only contains one risky lottery ℓ_r and the risk free lottery ℓ_S . With only one

risky lottery the state space is $\{\omega_1, \omega_2\}$, where ω_1 and payoff $\bar{\ell}_r$ occur with probability $p \in (0, 1)$, and ω_2 and payoff $\underline{\ell}_r$ occur with probability $1 - p$.

The risk-free lottery has the property that its outcome will always be known whether it was chosen or not. The same need not be true of the risky lottery. We let $q \in [0, 1]$ be the probability that an agent learns the outcome of the risky lottery, ℓ_r , conditional on choosing the safe lottery, ℓ_S . That is, conditional on choosing the risk-free lottery, with probability q the individual faces regret relevant partition $\pi_0 = \{F(S), F(\ell_r)\}$, and with probability $1 - q$ he faces partition $\pi'_S = \{F'(S)\}$ where $F'(S) = \Omega$. Utility is then given by

$$u^T(\ell_S, \omega, q) = \begin{cases} u(\ell_S) - q\kappa(u(\bar{\ell}_r) - u(\ell_S)), & \text{if } \omega = \omega_1 \\ u(\ell_S), & \text{if } \omega = \omega_2 \end{cases} \quad (5)$$

and

$$u^T(\ell_r, \omega, q) = \begin{cases} u(\bar{\ell}_r), & \text{if } \omega = \omega_1 \\ u(\underline{\ell}_r) - \kappa(u(\ell_S) - u(\underline{\ell}_r)), & \text{if } \omega = \omega_2 \end{cases} \quad (6)$$

Both the risky lottery, ℓ_r , and the risk-free lottery, ℓ_S , bring with them a benefit and a cost. The benefit is the direct choiceless utility associated with each; the cost is the psychological penalty that may be incurred in the event your choice is not optimal in the realised state. For the risky lottery, ℓ_r , both the expected benefit and expected cost are fixed. For the risk-free lottery, ℓ_S , the expected benefit is fixed but the expected cost is not. It may or may not be incurred.

Using (5) and (6), it is simple to calculate when the risky option is preferable to a regret averse individual. To make things as clear as possible, we normalise the choiceless utility function u so that $u(\underline{\ell}_r) = 0$ and $u(\ell_S) = 1$. With this, the condition for the risky lottery ℓ_r to be preferred to the risk-free lottery ℓ_S reduces to

$$u(\bar{\ell}_r) \geq \frac{1}{p} \left(\frac{1 + \kappa(1 - p(1 - q))}{1 + q\kappa} \right) \quad (7)$$

Expression (7) is bookended by two important cases. When $q = 1$ the decision maker will learn the outcome of the risky option no matter what. Here, there is no

distortion to the threshold rule relative to standard preferences, in that,

$$u(\bar{\ell}_r) \geq \frac{1}{p} \tag{8}$$

When $q = 0$ however, the decision maker knows that he will definitely *not* learn the outcome of the risky option unless he opts for it. Then (7) becomes,

$$u(\bar{\ell}_r) \geq \frac{1}{p}(1 + \kappa(1 - p)) \tag{9}$$

Given that $\kappa > 0$ and $p \in (0, 1)$, the inequality in (9) is a more demanding condition on the risky lottery ℓ_r than that in (8). That is, when a regret averse decision maker will certainly not find out the realisation of the risky option without choosing it, he requires it to have a more desirable payoff distribution. (It can be checked that the RHS of expression in (7) is strictly decreasing in q over the interval $[0, 1]$.)

The reason for the discrepancy above is that the risk-free lottery can provide insurance against regret. When $q = 0$, the risk-free lottery provides complete insurance against regret. There is complete asymmetry in anticipated regret: if the decision maker chooses the risky option, he knows he will be able to make an ex-post comparison and feel regret if the risky option is not successful, whereas if he chooses the risk-free lottery, he knows he will not be able to make such a comparison. Because of the insurance against regret offered by the risk-free lottery, the outcome of the risky lottery in the event of a success must be high enough to tempt the decision maker away from the security of the risk-free lottery. Ignorance is bliss. On the other hand, when $q = 1$, the risk-free lottery offers no insurance against regret. The decision maker knows he will learn the outcome of the risky lottery whether he chooses it or not. The regret considerations are symmetric and cancel each other out, and the individual's condition is the same as if he were regret neutral.

Intuitively, the benefit that the risky lottery must yield in order to be chosen becomes lower as q increases due to the reduction in anticipated regret. As q decreases (i.e., the likelihood of making an ex-post comparison in the case the risky option is not chosen reduces) the agent is increasingly - from an ex-ante perspective - insured against regret. And because of this insurance against potential regret, the risky lottery's outcome must increase to tempt the agent away from the safe option.

In the next subsection, we extend the setting to one with multiple regret averse agents, and we suppose that q is endogenously determined by the choices of others. In particular, we assume that q is increasing in the number of other agents who choose the risky lottery. This seemingly minor amendment turns a series of single-person decision problems into a multi-player behavioural game.

2.3 Strategic Setting: The Regret Game

2.3.1 The set up

We stick with the 2-state 2-lottery environment from the previous subsection. Suppose now that there are N decision makers, each of whom is regret averse, i.e., with preferences represented by u^T with $\kappa > 0$. The key modelling feature we introduce is that an individual's likelihood of learning about the risky outcome conditional on choosing the risk-free lottery is a function of the behaviour of others.

To capture the above, we define a symmetric N -player (simultaneous-move) behavioural game with common action set A and common utility function u^T .¹⁴ We identify the action set A with the 2-element choice set $\mathcal{L} = \{\ell_r, \ell_S\}$. Each player i , chooses an action $a_i \in A = \{\ell_r, \ell_S\}$, and has utility function $u_i^T : \mathbf{A} \times \{\omega_1, \omega_2\} \rightarrow \mathbb{R}$, where $\mathbf{A} := \prod_{j=1}^n A_j$, with typical element $\mathbf{a} = (a_1, \dots, a_n)$. From player i 's perspective, a pure action profile $\mathbf{a} \in \mathbf{A}$ can be viewed as (a_i, \mathbf{a}_{-i}) , so that $(\hat{a}_i, \mathbf{a}_{-i})$ will refer to the profile $(a_1, \dots, a_{i-1}, \hat{a}_i, a_{i+1}, \dots, a_n)$, i.e., the action profile \mathbf{a} with \hat{a}_i replacing a_i . The utility function of agent i is as defined in (5) and (6) save one difference: the probability that i learns the outcome of the risky lottery when it is not chosen, denoted q_i , depends on the behaviour of the other agents. Formally,

$$u_i^T((\ell_S, \mathbf{a}_{-i}), \omega) = \begin{cases} u(\ell_S) - q_i(\mathbf{a}_{-i})\kappa(u(\bar{\ell}_r) - u(\ell_S)), & \text{if } \omega = \omega_1 \\ u(\ell_S), & \text{if } \omega = \omega_2 \end{cases} \quad (10)$$

and

$$u_i^T((\ell_r, \mathbf{a}_{-i}), \omega) = \begin{cases} u(\bar{\ell}_r), & \text{if } \omega = \omega_1 \\ u(\underline{\ell}_r) - \kappa(u(\ell_S) - u(\underline{\ell}_r)), & \text{if } \omega = \omega_2 \end{cases} \quad (11)$$

where $q_i(\mathbf{a}_{-i})$ is assumed to be strictly increasing in the number of others that choose

¹⁴This is an abuse of notation as the domain on which utility is defined will now be extended to include the behaviour of others.

ℓ_r .¹⁵ That is, abusing notation somewhat, we let $|\mathbf{a}| = |(a_i, \mathbf{a}_{-i})| := \#\{j \neq i : a_j = \ell_r\}$. Then for any two profiles \mathbf{a}' and \mathbf{a}'' we have $q_i(|\mathbf{a}'|) > q_i(|\mathbf{a}''|)$ if and only if $|\mathbf{a}'_{-i}| > |\mathbf{a}''_{-i}|$. Finally, we assume that $q(0) = 0$ and $q(N - 1) = 1$.

We emphasise that, while player i 's utility depends on the choices of everyone, and hence the set up is a strategic game, the dependence is not direct. Rather it manifests through the likelihood that player j 's choice of lottery will impact player i 's ex-post regret relevant partition. It is individual j 's risk taking that can generate information for individual i , in turn altering i 's psychological payoffs, in turn altering the relative benefits of each choice. We now characterise the set of equilibria to this set up: showing how individual psychological motives can lead to socially interdependent decisions.

2.3.2 Equilibria

The above defines the Regret Game. There are three classes of the game, depending upon the parameters. When $u(\bar{\ell}_r) < 1/p$, it is a dominant choice for each player to choose the safe lottery ℓ_S . Similarly, when $u(\bar{\ell}_r) > \frac{1}{p}(1 + \kappa(1 - p))$, choosing the risky lottery ℓ_r is the dominant action. For intermediate values of $u(\bar{\ell}_r)$ however, the game is one of coordination and has two pure-strategy Nash equilibria: all players choose ℓ_S , and all players choose ℓ_r .¹⁶ Theorem 3 below states this formally.

Theorem 3. *In the regret game,*

1. *When $u(\bar{\ell}_r) < \frac{1}{p}$, uniform adoption of the risk-free lottery, ℓ_S , is the unique (dominant) pure strategy Nash equilibrium.*
2. *When $u(\bar{\ell}_r) > \frac{1}{p}(1 + \kappa(1 - p))$, uniform adoption of the risky lottery, ℓ_r , is the unique (dominant) pure strategy Nash equilibrium.*
3. *When $u(\bar{\ell}_r) \in [\frac{1}{p}, \frac{1}{p}(1 + \kappa(1 - p))]$, uniform adoption of the risk-free lottery, ℓ_S , is a pure strategy Nash equilibrium and uniform adoption of the risky lottery, ℓ_r , is a pure strategy Nash equilibrium*

¹⁵Assuming that q_i is increasing seems natural to us. The more individuals who choose the risky lottery, the harder it ought be not to learn about its performance ex-post. But while it seems implausible that the function q_i would be decreasing at any point, we leave its precise functional form unspecified. One can imagine settings in which q_i is linear, concave, and convex. One can even envisage settings where q_i is step function, in that the outcome of ℓ_r cannot be avoided once *enough* individuals have chosen it.

¹⁶There is also a completely mixed strategy equilibrium but we ignore it as it is very unstable.

Theorem 3 can be understood as follows. While the risky lottery, ℓ_r , brings a guaranteed utility, for every (other) individual who chooses the risky lottery, the expected utility associated with the risk-free lottery is decreasing. The reason for this being that the likelihood of learning about the alternative, and hence experiencing regret, is going up. Thus, for the parameters of case 3, we have a slightly unusual coordination game in that the number of others who choose ℓ_r is decreasing the associated net utility of choosing ℓ_S without improving the net utility of choosing ℓ_r .

2.3.3 Welfare and equilibrium selection

It is interesting to compare the (common) expected utility levels at each equilibrium in case 3. From (10) and (11) we can compute that uniform adoption of the risky lottery, ℓ_r , is Pareto optimal if and only if $u(\bar{\ell}_r) \geq \frac{1}{p}(1 + \kappa(1 - p))$. But this is precisely the threshold at which individuals would always choose the risky lottery in any case. Thus, for parameters such that the regret game is a coordination game (case 3), coordinating on the risk-free lottery is Pareto-optimal (and always preferred ex-ante).

There is a large literature addressing the tension that exists between the multiple equilibria in coordination games. The most commonly studied environment is a large-population binary-action game (e.g., the Stag Hunt) where the Pareto efficient equilibrium does not coincide with the *safe* risk-dominant equilibrium.¹⁷ Existing equilibrium selection techniques - be they evolutionary like *stochastic stability* (Kandori et al., 1993; Young, 1993) or higher-order belief-based like *global games* (Carlsson and Damme, 1993; Morris and Shin, 2003) - favour the equilibrium that is more difficult to destabilise (i.e., the risk-dominant one). For our regret game, simple algebra shows that $q \geq q^* = 1/2$ is the threshold at which the risky lottery becomes most desirable. But if we imagine a *steep* specification of q (in that $q_i(|\mathbf{a}|) \approx 1$ when $|\mathbf{a}| \geq 1$), then uniform adoption of the risky lottery would be the prediction. But that means the regret game has the quite curious property that the risk-free Pareto dominant outcome need not be risk-dominant (Harsanyi and Selten, 1988).

¹⁷Probably the best known experimental test of the efficiency vs safety issue is (Van Huyck et al., 1990) who study human behaviour in the *minimum effort game*, of which the Stag Hunt is a simple version.

2.3.4 Extensions

In what follows, we briefly describe some extensions of our model. Further details can be found in Appendix B.

In the regret game above we have assumed that everyone is regret averse. However, in reality it is possible that only some proportion of a population are regret averse ($\kappa > 0$) and the remainder are regret neutral ($\kappa = 0$). The first extension in Appendix B.1 considers the case in which some proportion of the population are similarly regret averse and the remaining proportion are regret neutral. The preferences over lotteries for regret neutral individuals are fixed no matter the information feedback. So these individuals always choose the lottery with the highest expected (choiceless) utility. When the environment becomes a game these individuals have a dominant strategy. The presence of regret neutral individuals makes it more difficult for the regret averse individuals to coordinate on the *other* lottery. The reason is that the regret neutral individuals will always choose one of the lotteries.

The second extension in Appendix B.1 considers the more general case where all individuals may have a different coefficient of regret aversion. Ex-ante, the regret coefficients are unknown, with everyone's coefficient being a random draw from some distribution over $[0, \infty)$. This generates a rich Bayesian game. In this sense, the first extension above is one of many possible realisations of the strategic environment.

Individuals who are not regret averse may be regret neutral or *rejoice lovers*. Rejoice is defined as the psychological gain that a decision maker – so called rejoice lover – experiences when the option chosen turns out to be better than the unchosen option. Appendix B.2 describes how rejoice can be incorporated in our model. When a rejoice lover does not learn the lottery's outcome, he will be *more* likely to choose the lottery, as he wants to know whether his choice was the best. It is easy to see that, while in equilibrium regret averse players will coordinate, rejoice lovers will anti-coordinate.

3 Experiments

We test the predictions of a two-player variant of the regret game through two experiments. Both experiments use a within-subject design and have two main goals. The first goal is identifying the participants who are regret averse. The second goal is testing whether regret averse participants behave as our model predicts. The two experiments

achieve these goals through two different designs, with the second experiment serving as a robustness check of the results by eliminating some confounds that may have been present in the first experiment.

3.1 Experiment 1

3.1.1 Design

Overview The experiment consists of two parts. In the first part we elicit participants' preferences over a riskless option (a sure amount of money) and a risky option (a lottery) under two different information environments. In the first environment participants learn the risky lottery's outcome even if they do not choose it, while in the second environment they do not learn the risky lottery's outcome unless chosen. These initial decisions allow us to classify participants into regret averse types and non regret averse types and to calibrate the parameters of the second part of the experiment. We also ask an additional question as a robustness check – to verify whether participants behave according to the type they were classified as. In the second part of the experiment, participants are matched in pairs and play the regret game. They must choose between a sure amount of money and a risky lottery. If they do not choose the lottery, they will learn its outcome only if their partner chose the lottery.

Part 1 In the first part of the experiment, Decisions 1 through 3, participants have to choose between a sure amount (€5 with certainty) and a risky lottery (€ x with 50% probability and €0 with 50% probability) under different conditions.

The standard incentive compatible mechanism for eliciting lottery thresholds is the BDM (Becker et al., 1964), and by Theorem 1 the BDM remains incentive compatible for regret averse individuals. We ask each participant to state the smallest lottery outcome x (henceforth *lottery threshold*) such that they prefer playing the lottery than receiving the sure amount. They can choose any number from the list $\{5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15\}$. After they submit their choice, the computer randomly picks a number from the same list, independently drawn for each participant. All the numbers are equally likely. If the number picked by the computer is smaller than the number x chosen by the participant, the sure amount is the implemented option, i.e., the participant receives €5. If the number picked by the computer is equal

or larger than the number x chosen by the participant, the lottery is the implemented option, i.e., the participant receives the number picked by the computer in € with 50% probability and €0 otherwise.¹⁸

The above is common to Decisions 1 and 2 (and, as discussed later, to Decision 3). The difference between the decisions is seen in the information feedback provided to the participants who do not choose the lottery. Labelling the number x chosen by participants in Decision 1 as x_1 and that chosen by participants in Decision 2 as x_2 , we have the following.

Decision 1. If a participant's choice, x_1 , results in the sure amount being the implemented option, they are nevertheless informed about the lottery's outcome.

Decision 2. If a participant's choice, x_2 , results in the sure amount being the implemented option, they are *not* informed about the lottery's outcome.

After each of the first two decisions, participants learn the number randomly picked by the computer, and thus whether the implemented option is the lottery or the sure amount. In the latter case, the information provided varies across the two decisions as described above.

Participants who choose a higher lottery threshold in Decision 2 (no information) than in Decision 1 (information), i.e., $x_2 > x_1$, are classified as *regret averse*. The sure amount is more appealing to them when feedback about the lottery is withheld, as it allows them to remain ignorant about the outcome of the unchosen option. Participants who choose the same or a lower lottery threshold in Decision 2 than in Decision 1, i.e., $x_2 \leq x_1$, are classified as *non regret averse*. These participants can be further classified in two categories. Participants who choose $x_2 = x_1$ are *regret neutral*, as they do not react to feedback about the unchosen option. Participants who choose $x_2 < x_1$ are *rejoice lovers*. The sure amount is less appealing to them when feedback about the lottery is withheld, as they cannot learn the lottery's outcome unless they choose it.¹⁹

To the best of our knowledge, our paper is the first to classify participants into regret averse and non regret averse types by using an incentive-compatible mechanism

¹⁸Whether the lottery is successful or not is perfectly correlated across participants.

¹⁹Since we use a within-subject design, we are potentially exposed to the risk of experimenter demand effect. However, we believe that it is very difficult for participants to figure out what the experimenter is trying to achieve through Decisions 1 and 2, and the heterogeneity in participants' responses seems to support our belief.

under different feedback conditions. However, the idea that the expectation of feedback affects anticipated regret and thus the behaviour of regret averse individuals has been documented by psychologists. For example, [Zeelenberg et al. \(1996\)](#) show that, when faced with the choice between two equally attractive gambles, most participants chose the gamble without feedback and thus ex-post comparisons of outcomes.²⁰

By classifying participants into types on the basis of a single decision, we are exposed to the risk of measurement error. However, if multiple decisions had been used to classify participants into types, we would have faced two other issues. First, earlier decisions may have affected later decisions, and varying the order to control for any possible order effect may not have been possible. Second, by going through multiple, similar questions participants may have felt bored and answered later question less accurately. In our second experiment, we use an alternative method to classify participants into types. While the new classification is still based on one decision only, by employing a different method it can serve as a robustness check for the classification into types used in our first experiment.

Finally, we ask participants an additional question (Decision 3) to verify whether they behave according to the type they were classified as through Decisions 1 and 2. Each participant is randomly paired with another participant and faces the same decision as in Decisions 1 and 2, with one difference: the feedback provided to the participants who does not play the lottery. If the number x chosen by a participant in Decision 3, labelled x_3 , results in the sure amount being the implemented option, they will be informed about the lottery's outcome only if their partner played the lottery. This means that if their partner played the lottery, participants are under the same information environment as in Decision 1, where they learn the lottery's outcome. If their partner did not play the lottery, they are under the same information environment as in Decision 2, where they do not learn the lottery's outcome unless they play it.

Participants are also asked what they believe is the lottery threshold x chosen by their partner in Decision 3. If their guess is within one of the number chosen by their partner, then they receive additional €1 at the end of the study (provided that Decision 3 is randomly selected for payment).

²⁰For a different method to classify subjects into regret averse and not, see [Bleichrodt et al. \(2010\)](#). Similarly to our design, [Imas et al. \(2021\)](#) compare participants' valuations for identical lotteries under two different feedback scenarios. However, they use a between-subject design and it is not possible to back out the number of regret averse individuals from their data.

As with Decisions 1 and 2, in Decision 3 participants are informed as to the number randomly picked by the computer, and thus whether the implemented option is the lottery or the sure amount.

Part 2 In the second part of the experiment (Decision 4 onwards), participants play the regret game described in Section 2.3. They must choose between a sure amount (earning €5 with certainty) and a lottery (earning “an amount” in € with 50% probability and €0 with 50% probability). Differently from the previous part of the experiment, the lottery’s outcome in the good state is given and, for each participant, it is determined using the lottery thresholds x_1 and x_2 elicited in Decisions 1 and 2.

The sure amount is known regardless of the decision made. If participants choose the sure amount, they learn the lottery’s outcome only if their partner chooses the lottery. This implies that if participants choose the lottery, they experience regret if the bad state materialises, whereas if they choose the sure amount, they experience regret if the good state materialises *and* they learn that.

This setting is identical in all the decisions in the second part of the experiment. However there are three type of decisions that differ in the amount the lottery yields in the good state corresponding to the three parts of the parameter space classified in Theorem 3. Two types of decisions are designed to test the regret game when it has a dominant strategy – in one the dominant strategy is the sure amount and in the other the dominant strategy is the lottery. The third type of decision is designed to test the regret game when it is a game of coordination. More interestingly, in each of these three decisions the amount that the lottery yields in the good state is individual-specific: it depends on a participant’s choices in Decisions 1 and 2. This ensures that each participant, despite potential heterogeneity in their (regret averse) preferences, plays the regret game in its three key cases.

Decision 4. The lottery’s outcome in the good state is *smaller* than the amount a participant chose in Decision 1, namely $x_1 - 2$. This allows us to test the prediction of the regret game when choosing the sure amount is the unique equilibrium in dominant strategies.

Decision 5. The lottery’s outcome in the good state is *bigger* than the amount a participant chose in Decision 2, namely $x_2 + 2$. This allows us to test the prediction of

the regret game when choosing the lottery is the unique dominant strategy equilibrium.

Decision 6. The lottery's outcome in the good state is *in between* the amount chosen in Decision 1 and that chosen in Decision 2, namely $\frac{x_1+x_2}{2}$. This allows us to test the prediction of the regret game when it is a game of coordination.

Decision 6, being the most interesting case of the regret game and thus the core part of the experiment, is repeated 20 times (Decisions 6 to 25).

Pairs are rematched before Decisions 4, 5 and 6. This means that a participant's partner could be the same as before or a different one. To mitigate potential dependence resulting from the repeated interaction of participants, we use matching groups of size 4. That is, for each participant, there are three potential participants who can be randomly assigned to them. In Decision 6 and its repetitions, participants keep the same partner.

From Decision 4 onwards, we elicit first order beliefs. We ask participants to guess whether their partner chose the sure amount of money or the risky lottery. If they guess their partner's choice, they receive additional €1 at the end of the study, if that decision is randomly selected for payment.

After the last iteration of Decision 6, i.e., Decision 25, the participants are asked some non-incentivised questions: the Regret Scale (Schwartz et al., 2002), the Big-Five personality traits (Gosling et al., 2003) and demographic characteristics.

Procedure The sessions were run in May 2017 in the experimental laboratory at the University of Bonn. The experiment was programmed and conducted with the software z-Tree (Fishbacher, 2007).

At the end of the experiment, each participant received: (i) a show-up fee of €4, (ii) the payment for one randomly chosen decision out of the 25 decisions made, (iii) the payment for one randomly chosen belief out of the 23 beliefs elicited, but only in case of a correct guess. On average, a participant earned €11.50.

3.1.2 Testable predictions

Decision 6 and its repetitions test Theorem 3 Part 3, i.e., they test the predictions of the regret game when it is a game of coordination with two pure strategy Nash Equilibria.

A regret averse type will choose the lottery if he believes that his partner chose the lottery, and the sure amount otherwise. This yields our first testable prediction.

Prediction 1. *In Decision 6 and its repetitions, believing that his partner will choose the lottery (sure amount) increases a regret averse agent's probability of choosing the lottery (sure amount).*

Decision 4 tests Theorem 3 Part 1, i.e., it tests the predictions of the regret game when choosing the sure amount is the unique equilibrium in dominant strategies. Decision 5 tests Theorem 3 Part 2, i.e., it tests the predictions of the regret game when choosing the lottery is the unique equilibrium in dominant strategies. When the regret game has a unique equilibrium in dominant strategies, an agent will choose it regardless of their regret preferences and regardless of their beliefs. This yields our second testable prediction.

Prediction 2. *Regret averse agents will choose the sure amount in Decision 4 and the lottery in Decision 5, regardless of their beliefs.*

We also have a third prediction, aimed at testing whether, under a partner-dependent information condition (basically a threshold-based variant of the regret game), participants behave according to the type they were classified as in Decisions 1 and 2. In Decision 3, regret averse (and rejoice loving) participants should choose the same lottery threshold as in Decision 1 if they believe that their partner played the lottery (i.e., if they expect to learn the lottery's outcome) and the same lottery as in Decision 2 if they believe that their partner did not play the lottery (i.e., if they expect not to learn the lottery's outcome). That is, they should choose a lottery threshold in between the lottery thresholds chosen under information and under no information.

As $x_2 > x_1$ for regret averse types and $x_2 < x_1$ for rejoice loving types, for regret averse (rejoice loving) participants x_3 should exceed (fall short of) x_1 and fall short of (exceed) x_2 . For regret neutral participants x_3 should equal x_1 and x_2 : they should choose the same lottery threshold as in Decisions 1 and 2, independently of their beliefs.

Prediction 3. *For regret averse (rejoice loving) agents, x_3 will exceed (fall short of) x_1 and fall short of (exceed) x_2 . For regret neutral agents, $x_1 = x_2 = x_3$.*

3.1.3 Results

A total of 144 subjects participated in the experiment with over 90% of them students and 56% of them female. The average age was 25 (24 among students and 27 among non students). In our sample, 22% of the participants chose $x_2 > x_1$ and are classified as regret averse. Half the participants chose $x_2 = x_1$ and are classified as regret neutral. The remaining participants chose $x_2 < x_1$ and are classified as rejoice loving. Figure 1 shows the distribution of the difference between x_2 and x_1 , which is a proxy for the strength of participants' regret aversion.




Figure 1: Difference between x_2 and x_1 .

To test Prediction 1, we estimate a logit panel-model with random effects, where the unit of observation is the individual observed in Decisions 6 through 25. We report the marginal effect of an agent's belief that his partner chose the lottery on her own likelihood of choosing the lottery.²¹ The dependent variable *lottery choice* takes value 1 if the agent chooses the lottery and 0 if he chooses the sure amount. We use the following independent variables. The variable *belief* equals 1 if the agent believes that his partner chose the lottery in the current round and 0 if he believes that his partner chose the sure amount. The variables *belief if regret averse* and *belief if non regret averse* equal *belief* if, respectively, the agent is regret averse and non regret averse (i.e.,

²¹We cluster the standard errors at the matching group-level.

regret neutral or rejoice lover). The variable *past regret* captures regret generated by previous decisions, and equals 1 (i) if an agent has not chosen the lottery in the previous round, while his partner has, and the lottery has been successful, or (ii) if an agent has chosen the lottery in the previous round and the lottery was not successful. It equals 0 otherwise.²² In column (1), we only control for *beliefs if regret averse* and *beliefs if non regret averse*. In column (2), we additionally control for *past regret*. In column (3), we additionally control for demographics (female dummy, student dummy and age).

We find that, consistently with Prediction 1, believing that their partner chose the lottery significantly increases regret averse participants' likelihood of choosing the lottery. We observe a positive impact of beliefs on choice also for non regret averse participants. However, for non regret averse participants the magnitude of the marginal effects is smaller than for regret averse participants.

Table 1: Impact of beliefs on choice (all rounds)

DV: lottery choice	(1)	(2)	(3)
belief if regret averse	0.175*** (0.032)	0.167*** (0.033)	0.166*** (0.033)
belief if non regret averse	0.172*** (0.022)	0.159*** (0.022)	0.157*** (0.022)
past regret	No	Yes	Yes
female dummy	No	No	Yes
student dummy	No	No	Yes
age	No	No	Yes
N	2880	2736	2736

Marginal effects from logit regression. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered at matching group-level. DV equals 1 if the agent chose the lottery and 0 otherwise.

Interestingly, when we focus on the last iteration of Decision 6 (Table 2), where we expect that learning may have helped participants to converge to equilibrium, we observe that the impact of beliefs on choice is larger than in Table 1 and highly significant for regret averse participants. In contrast, it is only marginally significant for non regret averse participants. Moreover, the difference in the magnitude of the marginal

²²Note that, while both (i) and (ii) can be interpreted as regret driven by past, unsuccessful decisions, their nature can potentially differ. While (i) captures peer-induced regret, as well as personal loss, (ii) only captures loss. Given that, we also repeat our regressions splitting *past regret* into two dummies respectively corresponding to cases (i) and (ii). Our results do not change.

effects between regret averse and non regret averse is larger in the last iteration than in all rounds pooled. Our first and core result follows.

Table 2: Impact of beliefs on choice (last round)

DV: lottery choice	(1)	(2)	(3)
belief if regret averse	0.301*** (0.091)	0.291*** (0.092)	0.284*** (0.093)
belief if non regret averse	0.195** (0.085)	0.179** (0.085)	0.162* (0.087)
past regret	No	Yes	Yes
female dummy	No	No	Yes
student dummy	No	No	Yes
age	No	No	Yes
N	144	144	144

Marginal effects from logit regression. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses, clustered at matching group-level. DV equals 1 if the agent chose the lottery and 0 otherwise.

Result 1. *In Decision 6 and its repetitions, believing that their partner chose the lottery (sure amount) significantly increases regret averse participants’ probability to choose the lottery (sure amount). This effect becomes stronger in the last iteration of the game.*

This result indicates that regret aversion drives coordination. However, in our setting there may have been two additional drivers of coordination: preferences for conformism (Charness et al., 2017, 2019) and inequity averse preferences (Fehr and Schmidt, 1999). As further discussed in Section 3.2, we ran a second experiment to eliminate these potential confounds.

Table 3 reports the percentages of participants choosing the lottery in Decisions 4 and 5, i.e., when the regret game has a dominant strategy. Consistently with Prediction 2, the large majority of regret averse participants chose the dominant strategy, i.e., the sure amount in Decision 4 and the lottery in Decision 5. As expected, also the large majority of non regret averse participants followed this pattern of decisions.²³

²³In Table 3 the non regret averse participants include regret neutral participants, and rejoice lovers when the game had a dominant strategy ($x_1 - x_2 \leq 2$). As expected, when the game had no dominant strategy ($x_1 - x_2 > 2$), the percentages of rejoice lovers choosing the sure amount in Decision 4 and the lottery in Decision 5 were remarkably lower.

Table 3: Participants choosing the lottery in Decisions 4 and 5

	regret averse	non regret averse
Agents who chose lottery in D4	6%	20%
Agents who chose lottery in D5	66%	80%

Table 4: Mean lottery thresholds chosen in Decisions 1-3 by type

	regret averse	regret neutral	rejoice lover	all
Decision 1 (x_1)	9.15 (2.40)	11.53 (2.04)	12.33 (1.98)	11.22 (2.39)
Decision 2 (x_2)	12 (2.56)	11.10 (2.70)	9.77 (2.34)	11.16 (2.40)
Decision 3 (x_3)	11.25 (2.59)	11.64 (2.30)	11.10 (2.70)	11.41 (2.47)
$x_1 - x_3$	-2.1***	-0.11	1.23**	-0.19
$x_2 - x_3$	0.75*	-0.54	-1.33*	-0.25
N	32	73	39	144

Standard deviation in parentheses. The Wilcoxon test tests $H_0 : x_1 - x_3 = 0$ and $H_0 : x_2 - x_3 = 0$.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Result 2. *Most regret averse agents choose the sure amount in Decision 4 and the lottery in Decision 5.*

Finally, as a robustness check, we verify whether under a partner-dependent information condition (Decision 3), participants behave consistently with the type they were classified as through Decisions 1 and 2. Table 4 presents the amounts chosen in Decision 1, Decision 2 and Decision 3 – overall and broken down by type. We find that, in line with Prediction 3, for regret averse participants the mean lottery threshold chosen in Decision 3, x_3 , is higher than mean x_1 and lower than mean x_2 . For rejoice loving participants, the opposite happens, and for regret neutral participants mean x_3 does not significantly differ from mean x_1 and mean x_2 .

To check whether these differences are statistically significant, we run the Wilcoxon equality test on matched data. The null hypotheses $H_0 : x_1 - x_3 = 0$ and $H_0 : x_2 - x_3 = 0$ are rejected for regret averse participants and rejoice loving participants, and not rejected for regret neutral participants. These results strongly support Prediction 3, thereby offering some reassurance that participants have been classified into types accurately.

Result 3. *The partner-dependent lottery threshold x_3 is significantly higher (lower)*

than the lottery threshold x_1 and significantly lower (higher) than the lottery threshold x_2 for regret averse (rejoice loving) participants. It is not significantly higher than x_1 and x_2 for regret neutral participants.

Prediction 3 implies that in Decision 3, the correlation between lottery threshold chosen and belief about the partner’s lottery threshold, denoted by χ , will be higher for regret averse participants than for non regret averse participants, i.e., $\chi_R > \chi_{NR}$. This is due to the fact that regret aversion induces a desire to coordinate. We find that χ_R equals 0.77 and χ_{NR} equals 0.47. The test for equality of correlation coefficients rejects the null hypothesis $H_0 : \chi_R = \chi_{NR}$. In particular, χ_R is significantly higher than χ_{NR} ($p = 0.01$).

3.2 Experiment 2

3.2.1 Design

Overview If participants have preferences for conformism (Charness et al., 2017, 2019), they may have chosen what they believed their partner to choose because in the game in Experiment 1 they did not know ex ante what the best choice was and thus preferred to do what their partner did in a previous round or what they believed their partner did in the current round. Moreover, if participants have inequity averse preferences (Fehr and Schmidt, 1999), they may have coordinated on the decision of their partner because in the game in Experiment 1 their earnings may have substantially differed from their partner’s earnings and they did not want to risk to earn less than their partner. To eliminate these potential confounds we designed and ran a second experiment.

The main goal of Experiment 2 is to test our model’s key prediction (i.e., that regret averse players try to coordinate with their partner) using a one shot variant of the regret game that eliminates the aforementioned potential confounds. The secondary goal of Experiment 2 is to provide an alternative and simple method to classify participants into regret averse and non regret averse types, which serves as a robustness check for the classification into types provided by Experiment 1.

Experiment 2 consists of two parts. In the first part, participants have to choose between a sure amount of money and a lottery. This decision is calibrated such that they prefer the sure amount. Then they are asked whether they want *to find out* how

much they would have earned had they chosen the risky lottery. If they choose *not to find out*, they forgo a small amount of money. This question allows us to classify participants into regret averse types and non regret averse types. In the second part of the experiment, participants are matched in pairs and play a variant of the regret game: they must choose whether they want to find out the risky lottery’s outcome or not, but they can avoid finding out only if their partner decides not to find out too.

Part 1 In this part subjects have to take 3 decisions.

Decision 1. We elicit each participant’s valuation of a lottery. We ask each participant for the smallest sure amount of money that they would choose over a risky lottery paying £80 with 20% probability and £0 with 20% probability.²⁴ We use an incentive compatible mechanism, the BDM, in the simpler version developed by [Healy \(2017\)](#). Participants are shown a list of 80 questions, each asking if they prefer the risky lottery (referred to as Option A) or a sure amount of money (referred to as Option B). The sure amount of money ranges from £1 (first question) to £80 (last question). Participants are then asked at which value of Option B they want to switch from Option A to Option B. Only 10% of the participants are paid for this decision.²⁵

Decision 2. We ask participants to choose between a sure amount of money (equal to the amount elicited in Decision 1 plus £2) and the risky lottery paying £80 with 20% probability and £0 with 80% probability. In order to be consistent with Decision 1, participants should choose the sure amount of money. All participants are paid for their decision in Decision 2.

Decision 3. All the participants who choose the sure amount in Decision 2 are asked whether they want to find out how much they would have earned, had they chosen the risky lottery.²⁶ By choosing not to find out, they forgo a small amount of money. In particular, if they chose to find out, they are informed about the lottery’s outcome at

²⁴To make the lottery easier to understand, we told them that the computer would randomly draw a ball from an urn containing 4 blue balls and 1 red ball. If the ball drawn was blue, they would earn £0; if the ball drawn was red, they would earn £80. We also showed a picture of the urn. See [Appendix C.2](#) for further details.

²⁵To ensure that only participants who read the instructions carefully remain in the study, participants who chose a number higher than twice the expected value of the lottery (24% of the initial sample) ended the experiment after this decision.

²⁶The participants who chose the lottery in Decision 2 (9%), thereby contradicting their previous decision, were then asked alternative questions and excluded from our data analysis. For further details, see [Appendix C.2](#).

the end of the experiment, and their earnings are increased by £0.04. If they chose not to find out, they are not informed about the lottery's outcome at the end of the experiment, and their current earnings are not increased by £0.04. Those who choose not to find out are classified as regret averse and those who chose to find out as non regret averse.

Part 2 In this part subjects play one shot of a variant of the regret game.

Decision 4. Participants are matched in pairs. Again, they must choose whether they want to find out the risky lottery's outcome or not. However, whether they learn the lottery's outcome also depends on their partner's behaviour. If they choose not to find out, they can avoid learning the lottery's outcome only if their partner also chose not to find out. Either this decision or the previous decision is randomly selected and implemented. If a participant chooses to find out in the randomly selected decision, their total earnings are increased by £0.04.

It is now easy to see how the design of Experiment 2 removes the potential confounds generated by preferences for conformism and inequity averse preferences. First, we use a one shot game in Decision 4, thus participants cannot imitate their partner's decisions in previous rounds. Second, in the variant of the game in Decision 4 participants can immediately identify the best decision to take given their preferences. This eliminates the concern generated by conformism motives, which would be at play if participants were unsure of the best thing to choose. Finally, in Decision 4 the potential earning difference generated by a participant's decision is negligible (at most £0.04). This eliminates the concern generated by inequity averse motives.

After making their decision, participants are also asked to guess their partner's decision. If they guess their partner's decision, they earn additional £0.50. Only one of the last two questions is randomly selected and implemented.

Differently from Experiment 1, in Experiment 2 the strategic decision aimed at testing the predictions of the regret game was not repeated. The reasons are the following. First, we wanted to avoid repeated-game effects and particularly the effect of conforming with the partner's previous decisions. Second, given that the strategic decision was simple and built on the previous decision, it did not appear necessary to provide participants with learning opportunities. Third, we thought that the effect of regret on behaviour may be more salient when the game is played only once.

Procedure Due to COVID 19-related lab closures, the sessions were run online in May 2021. To increase the robustness and external validity of our results, we used two different samples: students from Royal Holloway University of London and Prolific participants. The experiment was programmed and conducted with the software o-Tree (Chen et al., 2016).

At the end of the experiment, 10% of participants were paid for Part 1 and every participant was paid for Decision 2. If in the randomly selected decision out of Decision 3 and Decision 4, a participant chose to find out the lottery’s outcome, their total earnings were increased by £0.04. If they guessed their partner’s decision, they earned additional £0.50. On average, a participant earned £18.21.

3.2.2 Testable predictions

Decision 4 tests the prediction of the regret game when it is a game of coordination. In Decision 4, a regret averse participant will choose to find out the lottery’s outcome if he believes that his partner chose to find out the lottery’s outcome too. He will choose not to find out the lottery’s outcome if he believes that his partner also chose not to find out the lottery’s outcome. This implies that the share of regret averse agents choosing to find out under the belief that their partner chose to find out will be significantly higher than the share of regret averse agents choosing to find out under the belief that their partner chose not to find out. Similarly, the share of regret averse agents choosing not to find out under the belief that their partner also chose not to find out will be significantly higher than the share of regret averse agents choosing not to find out under the belief that their partner chose to find out. This yields our testable prediction.

Prediction 4. *The fraction of regret averse agents choosing the option that they believe their partner chose will be significantly higher than the fraction choosing the alternative option.*

3.2.3 Results

We have a sample of 213 participants who completed the experiment: 84 students from Royal Holloway University and 129 participants from Prolific. 54% of the participants were female, 44% were male and the remaining 2% classified themselves as “other”.

Table 5: regret averse subjects' choices and beliefs in Decision 4

	Believes partner finds out	Believes partner does not find out	Total
Finds out	5	0	5
Does not find out	8	12	20
Total	13	12	25

The gender distribution is very similar across our two subsamples. The average age was 27 (21 among Royal Holloway students and 32 among Prolific participants).

In our sample, 12% of the participants chose not to find out the lottery's outcome and are classified as *regret averse* and 88% of the participants chose to find out the lottery's outcome and are classified as *non regret averse*.

Table 5 shows the distribution of choices – between finding out and not finding out – and beliefs for regret averse participants. We can observe that 80% of the regret averse participants chose not to find out, as in the individual decision, and 20% chose to find out. It makes sense that either option gets chosen, as in Decision 4 there is no dominant strategy. All the participants who expect their partner to choose not to find out, also chose not to find out. Some participants chose not to find out even if they believed that their partner found out. This may have the following explanation. Given that we elicited point beliefs, participants who reported that their partner chose to find out could still believe that with some probability their partner chose not to find out. In this case it would be optimal not to find out, because the small amount of money to forgo is compensated by the reduction of the expected regret.

Table 6 shows the distribution of choices – between finding out and not finding out – and beliefs for non regret averse participants. Over 98% of the non regret averse participants (185 out of 188) chose to find out, as in the individual decision. This is expected, as finding out is a dominant strategy for them. It also shows that they are consistent across decisions. Out of these 185 participants choosing to find out, 10 believed that their partner chose not to find out. This is interesting, as it further confirms that finding out is a dominant strategy for them.

To check whether the differences observed in Table 5 are statistically significant, we run a t-test. Our null hypothesis is that the frequency with which that a regret averse participant chooses to find out under the belief that his partner did the same equals the frequency with which he chooses to find out under the belief that his partner chose

Table 6: Non regret averse subjects’ choices and beliefs in Decision 4

	Believes partner finds out	Believes partner does not find out	Total
Finds out	175	10	185
Does not find out	0	3	3
Total	175	13	188

not to find out. The relative frequencies are 1 and 0.61, respectively. We reject the null hypothesis ($p=0.0151$). Our results support Prediction 4.

Result 4. *The frequency with which regret averse participants choose the option that they believe their partner chose is significantly higher than the frequency with which they choose the alternative option.*

4 Conclusion

This paper began with the following simple observation: in many situations, ranging from technology adoption to ordering food in a restaurant, learning the outcome of unchosen alternatives is not guaranteed. Given that a decision maker can only regret a foregone alternative if she learns its outcome, what should be done? We showed how to incorporate this observation into the classic model of a decision maker who is regret averse. Our first contribution is a formalisation of ex post information structures that allow for the possibility that unchosen alternatives are / are not learned about. That is, the domain of preferences needs to be extended from simply ‘objects of choice’ to ‘objects of choice and their associated information environment’. For a given choice set, we provide a definition that ranks two informational environments according to which is “more informative”. And we show, in Theorem 2, that a more informative environment is never preferred for a regret-averse decision maker.

In Section 2.3, we allow for the possibility that the ex-post information environment that will be faced depends not only on one’s own choice but also on the choices of others. Thus, what on the surface appears to be a collection of individual decision problems - like for example ordering food in a restaurant - is in fact a rich multi-player behavioural game. The reason, of course, is that the decisions of others - one’s fellow diners - can be informative about foregone alternatives, and for a regret averse individual that matters.

We term this environment the regret game, and in Theorem 3 we classify conditions on preferences for which it is a coordination game with multiple equilibria.

We tested the predictions of our model through two experiments. Both experiments have two main goals: identifying the participants who are regret averse and testing whether they behave as our theory predicts. In the first experiment, we find that, as predicted by our model, regret averse participants try to coordinate with their partner. Believing that their partner chose an option significantly increases their likelihood of choosing that option. We observe a positive impact of beliefs on choice also for non regret averse participants. However, for non regret averse participants this impact is smaller. Moreover, when we focus on the last iteration of the game, the impact of beliefs on choice is larger and highly significant for regret averse participants, but is only marginally significant for non regret averse participants.

These results indicate that regret aversion drives coordination. However, preferences for conformism (Charness et al., 2017, 2019) and inequity averse preferences (Fehr and Schmidt, 1999) may have been two additional drivers of coordination. We ran a second experiment to eliminate these potential confounds.

The results of the second experiment support the key findings of the first experiment: regret averse participants try to coordinate with their partner. The probability that they choose what they think their partner chose is significantly higher than the probability that they choose the alternative option. Most notably, all the regret averse participants who believe that their partner chose to avoid information, choose to avoid information too. Instead, almost all of the non regret averse participants (over 98%) play the dominant strategy, and some of them do so under the belief that their partner chose the alternative option. Thus, once the aforementioned confounds are removed by design, we still find strong support for our model’s predictions.

It is traditional in economic modelling to view the information available to individuals as part of the environment and not something over which there is choice. While we have focused on regret, one can imagine a host of other behavioural explanations for why an individual might (i) want to avoid information, and then (ii) take steps in order to avoid information. It then immediately follows, that since the behaviour of others can impact one’s available information, there are strategic interdependencies at play. It is our hope that more models of the sort presented here can be developed to better-understand the role that evolving information plays for economic agents, and in

particular the way in which it connects them.

References

- Ariely, D., Levav, J., 2000. Sequential choice in group settings: Taking the road less traveled and less enjoyed. *Journal of Consumer Research* 27, 279–290.
- Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96. URL: <http://ideas.repec.org/a/eee/mateco/v1y1974i1p67-96.html>.
- Banerjee, A.V., 1992. A simple model of herd behavior. *The Quarterly Journal of Economics* 107, 797–817.
- Becker, G.M., DeGroot, M.H., Marschak, J., 1964. Measuring utility by a single response sequential method. *Behavioral Science* 9, 226–232.
- Bell, D.E., 1982. Regret in decision making under uncertainty. *Operations Research* 30, 961–981.
- Bénabou, R., 2013. Groupthink: Collective delusions in organizations and markets. *The Review of Economic Studies* 80, 429–462. URL: <http://www.jstor.org/stable/43551491>.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, pp. 992–1026.
- Bleichrodt, H., Cillo, A., Diecidue, E., 2010. A quantitative measurement of regret theory. *Management Science* 56, 161–175. URL: <http://www.jstor.org/stable/27784099>.
- Carlsson, H., Damme, E.v., 1993. Global games and equilibrium selection. *Econometrica* 61, 989–1018.
- Cesa-Bianchi, N., Lugosi, G., 2006. *Prediction, Learning, and Games*. Cambridge University Press. doi:[10.1017/CB09780511546921](https://doi.org/10.1017/CB09780511546921).
- Charness, G., Naef, M., Sontuoso, A., 2019. Opportunistic conformism. *Journal of Economic Theory* 180, 100–134.

- Charness, G., Rigotti, L., Rustichini, A., 2017. Social surplus determines cooperation rates in the one-shot prisoner's dilemma. *Games and Economic Behavior* 100, 113–124.
- Chen, D.L., Schonger, M., Wickens, C., 2016. otree – an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance* 9, 88–97.
- Cooper, D.J., Rege, M., 2011. Misery loves company: Social regret and social interaction effects in choices under risk and uncertainty. *Games and Economic Behavior* 73, 91–110. URL: <https://www.sciencedirect.com/science/article/pii/S0899825611000091>, doi:<https://doi.org/10.1016/j.geb.2010.12.012>.
- Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* 114, 817–868. URL: <http://www.jstor.org/stable/2586885>.
- Filiz-Ozbay, E., Ozbay, E.Y., 2007. Auctions with anticipated regret: Theory and experiment. *The American Economic Review* 97, 1407 – 1418.
- Fishbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Foster, D.P., Vohra, R.V., 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21, 40–55. URL: <https://www.sciencedirect.com/science/article/pii/S0899825697905959>, doi:<https://doi.org/10.1006/game.1997.0595>.
- Gali, J., 1994. Keeping up with the joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit and Banking* 26, 1–8.
- Gosling, S.D., Rentfrow, P.J., Jr, W.B.S., 2003. A very brief measure of the big-five personality domains. *Journal of Research in Personality* 37, 504–528.
- Harsanyi, J.C., Selten, R., 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press.

- Hart, S., Mas-Colell, A., 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 1127–1150. URL: <http://www.jstor.org/stable/2999445>.
- Healy, P.J., 2017. Epistemic experiments: Utilities, beliefs, and irrational play.
- Imas, A., Lame, D., Wilson, A.J., 2021. Reversals between one-shot and repeated decisions in incentive design: the case of regret.
- Kandori, M., Mailath, G.J., Rob, R., 1993. Learning, mutation, and long run equilibria in games. *Econometrica* 61, 29–56.
- Kreps, D.M., Porteus, E.L., 1978. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46, 185–200. URL: <http://www.jstor.org/stable/1913656>.
- Loomes, G., Sugden, R., 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* 92, 805–824.
- Morris, S., Shin, H.S., 2003. Global games: Theory and applications, in: in “Advances in Economics and Econometrics, the Eighth World Congress”, Dewatripont, Hansen and Turnovsky, Eds.
- Sarver, T., 2008. Anticipating regret: Why fewer options may be better. *Econometrica* 76, 263–305. doi:[10.1111/j.1468-0262.2008.00834.x](https://doi.org/10.1111/j.1468-0262.2008.00834.x).
- Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., Lehman, D.R., 2002. Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology* 83, 1178–1197.
- Van Huyck, J.B., Battalio, R.C., Beil, R.O., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review* 80, 234–48. URL: <http://ideas.repec.org/a/aea/aecrev/v80y1990i1p234-48.html>.
- Young, H.P., 1993. The evolution of conventions. *Econometrica* 61, 57–84.
- Zeelenberg, M., Beattie, J., van der Pligt, J., de Vries, N.K., 1996. Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes* 65, 148–158.

Appendix

A Proofs

Proof of Theorem 1

Proof.

$$\begin{aligned}
\mathbb{E}\left[u^T(\ell_i(\omega))\right] &\geq \mathbb{E}\left[u^T(\ell_j(\omega))\right] \\
\iff \mathbb{E}\left[u(\ell_i(\omega))\right] - \mathbb{E}[R(\ell_i; \omega)] &\geq \mathbb{E}\left[u(\ell_j(\omega))\right] - \mathbb{E}[R(\ell_j; \omega)] \\
\iff \mathbb{E}\left[u(\ell_i(\omega))\right] - \kappa \mathbb{E}\left[\max_{\ell \in \mathcal{L}} (u(\ell(\omega)) - (u(\ell_i(\omega))))\right] &\geq \mathbb{E}\left[u(\ell_j(\omega))\right] - \kappa \mathbb{E}\left[\max_{\ell \in \mathcal{L}} (u(\ell(\omega)) - (u(\ell_j(\omega))))\right] \\
\iff (1 + \kappa) \mathbb{E}\left[u(\ell_i(\omega))\right] - \mathbb{E}\left[\max_{\ell \in \mathcal{L}} (u(\ell(\omega)))\right] &\geq (1 + \kappa) \mathbb{E}\left[u(\ell_j(\omega))\right] - \mathbb{E}\left[\max_{\ell \in \mathcal{L}} (u(\ell(\omega)))\right] \\
\iff (1 + \kappa) \mathbb{E}\left[u(\ell_i(\omega))\right] &\geq (1 + \kappa) \mathbb{E}\left[u(\ell_j(\omega))\right] \\
\iff \mathbb{E}\left[u(\ell_i(\omega))\right] &\geq \mathbb{E}\left[u(\ell_j(\omega))\right]
\end{aligned}$$

□

Proof of Theorem 2

Proof. Clearly we can ignore lotteries ℓ_i for which $\pi'_i = \pi''_i$. So consider a lottery k such that π'_k is strictly finer than π''_k (note that Definition 1 requires that there is some such lottery). By the definition of more informative there is a non-empty event F that is an element of π''_k but not π'_k , and moreover, F is the union of disjoint elements of π'_k . Let \mathcal{O}'_k be the set of lotteries that are learned about in environment Π' and let \mathcal{O}''_k be the set of lotteries that are learned about in environment Π'' . Clearly, $\mathcal{O}''_k \subset \mathcal{O}'_k$, and so it must be that for every state ω we have

$$\max_{\ell_h \in \mathcal{O}'_k} \ell_h(\omega) \geq \max_{\ell_h \in \mathcal{O}''_k} \ell_h(\omega),$$

with the inequality being strict for at least one state.

Since choiceless utility does not depend on the informational environment, we need only consider the regret. Now note that the inequality above can then be adapted to

show that for all states $\omega \in \Omega$,

$$\kappa \left(\max_{\ell_h \in \mathcal{O}'_k} u(\ell_h(\omega)) - u(\ell(\omega)) \right) \geq \kappa \left(\max_{\ell_h \in \mathcal{O}''_k} u(\ell_h(\omega)) - u(\ell(\omega)) \right)$$

with the inequality being strict for at least one state.

Thus, for every lottery the total utility in a given state is no higher under information environment Π' than under Π'' and for some lottery there is a state at which the total utility is strictly lower. The result then follows. \square

Proof of Theorem 3

Proof. Parts 1 and 2 are immediate as players have a dominant strategy over each range of parameters.

Consider part 3. It is easy to see that both symmetric profiles are Nash Equilibria over this range. To see that the symmetric outcomes are the only pure strategy Nash Equilibria, suppose to the contrary. That is, suppose there is a pure strategy Nash equilibrium, $\hat{\mathbf{a}}$, in which some individual, say i , chooses ℓ_S and another individual, say j , chooses the risky lottery, ℓ_r . Since q is defined as a strictly increasing function from $\{0, 1, 2, \dots, N-1\}$ to $[0, 1]$, we have that $q_i(\hat{\mathbf{a}}) = q(\#\{k \neq i : \hat{a}_k = \ell_r\}) > q(\#\{k \neq j : \hat{a}_k = \ell_r\}) = q_j(\hat{\mathbf{a}})$. But this contradicts the fact that ℓ_S is optimal for individual i and ℓ_r is optimal for individual j , and so profile $\hat{\mathbf{a}}$ cannot be a pure strategy Nash equilibrium. \square

B Extensions

B.1 Heterogeneous regret preferences

There are many ways to extend our framework to incorporate heterogeneity in levels of regret aversion. Below we consider two such avenues. The first is a special case of the second.

A simple example with heterogeneity in levels of regret aversion. Suppose the population is comprised of two distinct groups, with all individuals in a given group sharing a common coefficient of regret aversion. For the sake of simplicity we assume that those in the first group have coefficient of regret aversion $\kappa > 0$ while those in the

second group have coefficient of regret aversion equal to zero. We refer to those in the first group as regret averse and those in the second group as regret neutral.

We assume that the parameters of the problem are given by condition 3 of the Theorem, i.e., that $u(\bar{\ell}_r) \in [\frac{1}{p}, \frac{1}{p}(1 + \kappa(1 - p))]$. For these parameters, the optimal choice for the regret averse individuals depends on population behaviour.²⁷

With the population size given by N , suppose that the number of regret averse individuals is m (and hence the number of regret neutral individuals is $N - m$). Without loss of generality we can relabel everyone such that the regret averse individuals are listed first (i.e., 1 through m). This set up is a variant of the regret game, in which the strategy chosen by the $N - m$ regret neutral individuals will be ℓ_r since this is the dominant strategy for these individuals.

Given that the regret neutral individuals have a dominant strategy, there are two candidates for pure strategy equilibria, \mathbf{a}^* and \mathbf{a}^{**} , where,

$$\mathbf{a}^* = \left(\underbrace{\ell_S, \dots, \ell_S}_{\text{behaviour of } m \text{ regret averse}}, \quad \underbrace{\ell_r, \dots, \ell_r}_{\text{behaviour of } N-m \text{ regret neutral}} \right)$$

$$\mathbf{a}^{**} = \left(\underbrace{\ell_r, \dots, \ell_r}_{\text{behaviour of } m \text{ regret averse}}, \quad \underbrace{\ell_r, \dots, \ell_r}_{\text{behaviour of } N-m \text{ regret neutral}} \right)$$

Clearly \mathbf{a}^{**} is always a pure strategy equilibrium. The issue of interest is under what conditions on group size m will \mathbf{a}^* be a pure strategy equilibrium. The issue for regret averse individuals now is that since $q(N - m) > 0$, there is always a positive probability that the outcome of the risky lottery will be learned ex post. Thus the risk free lottery no longer provides full insurance from regret for the regret averse individuals, and the relative payoff advantage to them from coordinating on the risk free lottery is diminished. In fact, if m is ‘too low’, then it is possible that there are simply ‘too many’ individuals who have ℓ_r as a dominant strategy for the m regret averse individuals to coordinate on the risk free lottery. The precise value of m for which this happens will depend upon the functional form of q_i but it can easily be calculated.

A rich example with heterogeneity in levels of regret aversion. In the above example, we fixed things such that there were only two levels of regret aversion, κ and 0.

²⁷If we did not make this assumption, then the regret averse individuals also have a dominant strategy (meaning that everybody has a dominant strategy so the situation is ‘obvious’).

Here we show how to model a Bayesian game, of which the above complete information example was only one of an infinite number of realisations of the strategic environment.

By this we mean the following. Suppose that everyone's regret coefficient is unknown ex-ante, but rather each individual i will have regret coefficient κ_i that is the realisation of a random variable drawn from some distribution with support $[0, \infty)$. Thus, the realised game will be one in which the individuals can be ordered 1 through N according to their realised coefficient of regret aversion (let the ordering be such that $\kappa_i < \kappa_j$ whenever $1 \leq i < j \leq N$).

It is clear that any individual with realised coefficient of regret aversion $\kappa < \frac{u(\bar{\ell}_r) - 1}{1 - p}$ will always choose the risky lottery since it is dominant to do so. With a slight abuse of terminology we refer to these individuals as regret neutral as, while they may have a coefficient of regret aversion that differs from 0, their coefficient of regret aversion is sufficiently low that it is a dominant strategy for them to choose the risky lottery. For all the remaining individuals, there is a coordination game being played that has the flavour of the complete information example above. That is, they are playing a coordination game amongst themselves, but there is an outside group of individuals who make choosing the risky lottery relatively more attractive. Of course, this happens for every possible realisation of the random vector of coefficients of regret aversion, $(\kappa_1, \dots, \kappa_N)$. Given knowledge of the distribution that determines this random vector, the Bayesian Nash equilibria to this game can be computed.

B.2 Rejoice

This subsection describes how our model can be extended to incorporate rejoicing into it. In the two state world, the utility of a regret individual is given by

$$U^R(\ell_S, \omega, q) = \begin{cases} u(\ell_S) - q\kappa(u(\bar{\ell}_r) - u(\ell_S)), & \text{if } \omega = \omega_1 \\ u(\ell_S), & \text{if } \omega = \omega_2 \end{cases} \quad (\text{A1})$$

and

$$U^R(\ell_r, \omega, q) = \begin{cases} u(\bar{\ell}_r), & \text{if } \omega = \omega_1 \\ u(\underline{\ell}_r) - \kappa(u(\ell_S) - u(\underline{\ell}_r)), & \text{if } \omega = \omega_2 \end{cases} \quad (\text{A2})$$

Now suppose that in addition to experiencing regret, individuals also experience rejoice when their choice turned out ex-post optimal. This is captured by amending equations (A1) and (A2) above to include a utility benefit, evaluated as a function

of the difference between ‘choiceless utility from chosen lottery’ and ‘choiceless utility from lottery not chosen’, that enters positively when the choice made turns out optimal ex-post. Formally, the equations above are replaced by

$$U^R(\ell_S, \omega, q) = \begin{cases} u(\ell_S) - q\kappa_1(u(\bar{\ell}_r) - u(\ell_S)), & \text{if } \omega = \omega_1 \\ u(\ell_S) + q\kappa_2(u(\ell_S) - u(\underline{\ell}_r)), & \text{if } \omega = \omega_2 \end{cases} \quad (\text{A3})$$

and

$$U^R(\ell_r, \omega, q) = \begin{cases} u(\bar{\ell}_r) + \kappa_2(u(\bar{\ell}_r) - u(\ell_S)), & \text{if } \omega = \omega_1 \\ u(\underline{\ell}_r) - \kappa_1(u(\ell_S) - u(\underline{\ell}_r)), & \text{if } \omega = \omega_2 \end{cases} \quad (\text{A4})$$

where κ_1 is the coefficient of regret aversion (replacing the κ from equations (A1) and (A2) and κ_2 weights the rejoice experienced. We note that there is again a q on the rejoice term associated with choosing the safe lottery, ℓ_S . As with the discussion in the main text, this is because the outcome of the risky lottery is not learned with certainty when the safe lottery is chosen. In this case however, the lack of certainty associated with risky lottery harms an individual prone to rejoice, as the fact that choosing the safe lottery turned out optimal may never be learned.

C Experimental instructions

C.1 Experiment 1

Welcome to the study.

Please note that you may not talk to the other participants at any time during the entire study. Should this happen, we will be forced to terminate the study.

Please read these instructions carefully.

In this study, you will make **25 decisions**.

At the end of the study, you will be paid in cash for **ONE** of these 25 decisions, picked at random by the computer.

Each decision is equally likely to be picked. So you should regard each decision as if it was the relevant one.

In addition, we will ask you **23 additional questions**.

At the end of the study, ONE of the additional questions will be randomly picked and paid on the base of your answer.

Each question is equally likely to be picked. So you should answer each question, as if it were the relevant one.

In addition, you will receive €4 for your participation in the study.

Instructions about the first part of the study

Please read these instructions carefully.

In the first part of the study, we will give you two options.

Left	Right
€x with 50% probability and €0 with 50% probability	€5 with certainty

First, you must specify **the smallest number x, such that you would prefer option “Left” to option “Right”**. You can choose any number in the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

After you have submitted your decision, the computer will randomly pick a number from the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

All numbers have the same probability of being picked.

If the number picked by the computer is **smaller** than the number x you chose, then **Right** will be the selected option. This means that you will get €5.

If the number picked by the computer is **equal to or bigger** than the number x you chose, then **Left** will be the selected option. This means that you will get the number of € picked by the computer with 50% probability and 0 with 50% probability.

If you have any questions, please raise your hand and we will come to you.

[Control questions]

Decision 1 (on screen only)

You have two options.

Left	Right
€x with 50% probability and €0 with 50% probability	€5 with certainty

First, you must specify the **smallest** number x, such that you would prefer option “Left” to option “Right”. You can choose any number in the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

After you have submitted your decision, the computer will randomly pick a number from the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

All numbers have the same probability of being picked.

If the number picked by the computer is **smaller** than the number x you chose, then **Right** will be the selected option. This means that you will get €5.

If the number picked by the computer is **equal to or bigger** than the number x you chose, then **Left** will be the selected option. This means that you will get the number of € picked by the computer with 50% probability and 0 with 50% probability.

Important information

After you have submitted your decision, the computer **will let you know** the outcome of option “Left” even if you have chosen a number x such that option “Right” is selected.

This means that, if “Right” is the selected option, you will learn nevertheless how much you would have earned, had “Left” been the selected option.

Decision 1. I prefer option Left to option Right if x is at least equal to _____.

Decision 2 (on screen only)

You face the same decision as before.

Left	Right
€ x with 50% probability and €0 with 50% probability	€5 with certainty

First, you must specify the **smallest** number x , such that you would prefer option “Left” to option “Right”. You can choose any number in the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

After you have submitted your decision, the computer will randomly pick a number from the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

All numbers have the same probability of being picked.

If the number picked by the computer is smaller than the number x you chose, then Right will be the selected option. This means that you will get €5.

If the number picked by the computer is equal to or bigger than the number x you chose, then Left will be the selected option. This means that you will get the number of € picked by the computer with 50% probability and 0 with 50% probability.

The only difference is the following. After you have submitted your decision, the computer **will NOT let you know** the outcome of option “Left” if you have chosen a number x such that option “Right” is selected.

This means that, if “Right” is the selected option, you will NOT learn how much you would have earned, had “Left” been the selected option.

Decision 2. I prefer option Left to option Right if x is at least equal to _____.

Instructions about decisions 3, 4, 5 and 6

Please read these instructions carefully.

Before each of the next four decisions, you will be randomly assigned to another participant.

In this lab, there are 3 potential participants who can be randomly assigned to you.

At no time will you find out the identity of the other participant.

If you have any questions, please raise your hand and we will come to you.

Decision 3 (on screen only)

You are randomly paired with another participant.

Each of you faces the same decision as before.

First, you must specify the **smallest** number x , such that you would prefer option “Left” to option “Right”. You can choose any number in the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

After you have submitted your decision, the computer will randomly pick a number from the list (5; 6; 7; 8; 9; 10; 11; 12; 13; 14; 15).

All numbers have the same probability of being picked.

If the number picked by the computer is **smaller** than the number x you chose, then **Right** will be the selected option. This means that you will get €5.

If the number picked by the computer is **equal to or bigger** than the number x you chose, then **Left** will be the selected option. This means that you will get the number of € picked by the computer with 50% probability and 0 with 50% probability.

The only difference is the following. If you have chosen a number x such that option “Left” is selected, you always know the outcome of option “Left”.

If you have chosen a number x such that option “Right” is selected, the computer **will inform you about the outcome of option “Left” only if “Left” is the selected option for your partner**. This means that, if “Left” is the selected option for your partner, you will learn nevertheless (as in Decision 1) how much you would have earned, had “Left” been the selected option. If “Right” is the selected option for your partner, you will NOT learn (as in Decision 2) how much you would have earned, had “Left” been the selected option.

Decision 3. I prefer option Left to option Right if x is at least equal to _____.

Question about partner’s choice

Now we ask you about your partner’s choice.

Which number do you think that your partner chose? _____.

If you guess the exact number chosen by your partner, or a number 1 point lower or 1 point higher than the number chosen by your partner, you will be paid €1 more at the end of the study.

Instructions about the second part of the study

Please read these instructions carefully.

In the second part of the study, we will give you two options.

Left	Right
€“a number” with 50% probability and €0 with 50% probability	€5 with certainty

Unlike in the previous part of the study, we will now give you a number.

So instead of choosing a number, you will choose either Left or Right.

If you have chosen option “Left”, you will **always** know the outcome of option “Left”.

If you have chosen option “Right”, **the computer will inform you about the outcome of option “Right” only if your partner has chosen “Right”**.

This means that if you have selected option “Right” and your partner option “Left”, you still learn (as in decision 1) how much you would have earned if you had chosen option “Left”. If you and your partner have chosen “Right”, you will NOT learn (as in decision 2) how much you would have earned if you had chosen “Left”.

If you have any questions, please raise your hand and we will come to you.

Decision 4 (on screen only)

You are now paired with another participant.

This participant could be either the same as before or a different one.

Left	Right
€ x chosen in D1-2 with 50% probability and €0 with 50% probability	€5 with certainty

If you have chosen option “Left”, you will **always** know the outcome of option “Left”.

If you have chosen option “Right”, **the computer will inform you about the outcome of option “Left” only if your partner has chosen option “Left”**.

This means that if you have chosen option “Right” and your partner option “Left”, you will learn nevertheless (as in Decision 1) how much you would have earned, had you chosen option “Left”. If both you and your partner have chosen option “Right”, you will NOT learn (as in Decision 2) how much you would have earned, had you chosen option “Left”.

Decision 4. Do you prefer option Left or option Right? [Right/Left]

Question about partner’s choice

Do you think that your partner chose option Left or option Right? [Right/Left]

If you guess this right, you will be paid €1 more at the end of the study.

Decision 5 and decision 6 are exactly the same as decision 4, but the outcome of the lottery in case of success is, respectively, the x chosen in decision 2 plus 2, and the sum of the x chosen in decision 1 and the x chosen in decision 2, divided by 2.

Instructions about the third part of the study

Please read these instructions carefully.

In the third part of the study, **you will repeat decision 6 for other nineteen times.**

In these additional nineteen decisions, you will still be paired with the **same partner.**

If you have chosen option “Left”, you will **always** know the outcome of option “Left”.

If you have chosen option “Right”, **the computer will inform you about the outcome of option “Right” only if your partner has chosen “Right”.**

This means that if you have selected option “Right” and your partner option “Left”, you still learn (as in decision 1) how much you would have earned if you had chosen option “Left”. If you and your partner have chosen “Right”, you will NOT learn (as in decision 2) how much you would have earned if you had chosen “Left”.

If you have any questions, please raise your hand and we will come to you.

C.2 Experiment 2

[Consent form]

Thank you for being a participant in this study. The study will take approximately 15 minutes.

Please give this study your full attention. You will have a limited amount of time to complete the study. Every screen has a timer. If time runs out, you will be excluded from the study and receive no payment.

As mentioned in the invitation email and in the reminder email, participants in this study will be paid via PayPal. Hence having a PayPal account and providing us with your PayPal email address is a prerequisite for participating in this study.

Please, enter your PayPal email address here: _____.

In the sessions run on Prolific, it was not necessary to ask for participants’ PayPal email address for payment, so the PayPal-related sentences were removed from the instructions.

Part 1

In this part of the study, you will be asked a list of questions.


10% of the participants in this study will be paid for this part of the study.

Part 1

I am going to ask you the following list of questions.

In each question, there are two options. Option A is always **drawing lots**. Option B is always **an amount of money**.

Drawing lots means the following. There are 5 balls in a jar, 4 are blue and 1 is red. The computer will randomly choose one ball from the jar. If the ball drawn is blue, you will earn £0. If the ball drawn is red, you will earn £80.



This means that Option A gives a 20% chance of earning £80 and an 80% chance of earning £0. Note that the expected value of Option A is £16. This means that on average Option A pays £16.

[table]

In each question you choose either Option A (drawing lots) or Option B (an amount of money). After you answer all 80 questions, I will randomly pick **one** question and pay you for your decision in that question. Each question is equally likely to be chosen for payment.

Obviously you have no incentive to lie on any question, because if that question gets chosen for payment then you would end up with the option you like less.

I assume you are going to choose Option A in at least the first few questions, but at some point switch to choosing Option B. So, to save time, just tell me at which value of Option B you would switch. I can then “fill out” your answers to all 80 questions based on your switch decision (choosing Option A for all questions before your switch decision, and Option B for all questions at and after your switch decision).

I will still draw one question randomly for payment. Again, if you lie about your true switch decision you may end up getting paid an option you like less.

At which value of Option B do you want to switch from Option A to Option B? Remember that Option A pays £0 with 80% probability and £80 with 20% probability. [*dropdown menu with all integers from 1 to 80*]

We will give you [*countdown timer*] minutes before you can submit your answer, so you can think about it carefully. After [*countdown timer*] minutes, the submit button below will become active.

Part 2

In this part of the study, you will be asked three questions: Questions 1, 2 and 3.

Every participant will be paid for their decision in Question 1.

Either Question 2 or Question 3 will be randomly selected and implemented.

Question 1

In Part 1 you switched at £[*amount chosen in previous question*].

In this question there are two options: an amount of money and drawing lots. Remember that drawing lots means that the computer will draw a ball from a jar containing 4 blue balls and 1 red ball. If the ball drawn is blue, you will earn £0. If the ball drawn is red, you will earn £80.

Which of these two options do you prefer?

- *amount chosen in previous question + 2*
- Drawing lots

Remember that every participant will be paid for their decision in this question.

Question 2

In Question 1 you chose £[*amount chosen in previous question+2*] over drawing lots.

Just so you know, the computer drew lots anyways. So it is possible to find out whether you would have earned £80 or £0 if you had chosen drawing lots in Question 1.

Which of these two options do you prefer?

- Not find out
- Find out

If you choose **Not find out**, you will avoid knowing whether you would have earned £80 or £0. Then your current earnings would be £[*amount chosen in previous question+2*].

If you choose **Find out**, you will not avoid knowing whether you would have earned £80 or £0, but £0.04 will be added to your earnings from Question 1. Then your current earnings would be £[*amount chosen in previous question*+2+0.04].

Remember that either Question 2 or Question 3 will be randomly selected and implemented.

Question 3

You are now paired with another participant in this study. This participant is now your partner. Your partner was asked the same questions as you up to now.

In Question 1 you chose £[*amount chosen in previous question*+2] over drawing lots.

However, recall than, as mentioned in Question 2, the computer drew lots anyways. So it is possible to find out whether each of you would have earned £80 or £0 if you had chosen drawing lots in Question 1.

If you choose **Not find out**, you will avoid knowing whether you would have earned £80 or £0 only if your partner also chose **Not find out**. Then your current earnings would be £[*amount chosen in previous question*+2].

If you choose **Find out**, you and your partner will not avoid knowing whether each of you would have earned £80 or £0, but £0.04 will be added to your earnings from Question 1. Then your current earnings would be £[*amount chosen in previous question*+2+0.04].

After making your choice, you will be asked to guess your partner's choice. You will earn additional £0.50 if you correctly guess your partner's choice.

Which of these two options do you prefer?

- Not find out
- Find out

Which option do you think your partner chose?

- My partner chose: Not find out
- My partner chose: Find out

Remember that either Question 2 or Question 3 will be randomly selected and implemented.

After this, there was one screen with basic demographic questions (age and gender) and one screen with results and earnings.

Participants who chose more than twice the expected value of the lottery in Part 1 ended the experiment after that decision. Participants who chose the risky lottery in Question 1, thereby exhibiting inconsistent preferences, were asked two alternative questions instead of Questions 2 and 3 (two individual questions involving a choice between a sure amount and a lottery).