University of Cyprus
Department of Economics

Working Paper 03-2025

Failures of Contingent Thinking and the Winner's Curse

Philippos Louis

Department of Economics, University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus
www.ucy.ac.cy/econ/?lang=en

Failures of Contingent Thinking and the Winner's Curse.

Philippos Louis[*]

October 28, 2025

Abstract

I design a within-subject experiment to investigate why individuals fall victim to the winner's curse. A known explanation is a failure of contingent thinking (FCT). My design disentangles the effects of pure FCT from cursedness—the failure to recognize the correlation between others' information and actions. Results show that many participants exhibit FCT without being cursed, while a similar fraction display cursed reasoning. Only a minority avoid both errors. By estimating structural models of cognitive reasoning, including one of pure FCT, I provide further support for these findings, clarifying distinct cognitive mechanisms underlying suboptimal behavior.

Keywords: *winner's curse, contingent thinking, correlation neglect, cursedness, lab experiment, within-subject design.*

[*]Department of Economics, Univesity of Cyprus, Nicosia, Cyprus. louis.philippos@ucy.ac.cy.
I gratefully acknowledge financial support of the Spanish Ministry of Science and Innovation and FEDER through grant "Consolidated Group-C" ECO2008-04756, as well as the Swiss National Science Foundation (SNSF 135135) and the European Research Council (ERC Advanced Investigator Grant, ESEI-249433). I wish to thank Joan de Martí, Salvador Barbera, Antonio Cabrales, Caterina Calsamiglia, Vincent Crawford, Jacob Goeree, Christos Ioannou, Nagore Iriberri, Pedro Rey Biel, Joel Sobel, John Wooders, Emanuel Vespa and especially Jingjing Zhang for helpful comments and suggestions. Also I thank Riste Gjorgjiev, Sabine Flamand, Orestis Troumpounis and Martin Schütz for help in conducting the experiments.
This paper is partially based on work included in a chapter of my thesis circulated under the tittle: *"Seeing is believing: an experiment in strategic thinking"*, and using a different set of experiments.
The experiment reported here is not preregistered. When it was first designed and run (2012) preregistration was not yet established as a norm for economic lab experiments. I try to be clear in the main text in distinguishing confirmatory from exploratory results.
Previous versions of this manuscript that were circulated and cited under the title "The barrel of apples game: Contingent thinking, learning from observed actions, and strategic heterogeneity" are superseded by the current version.

1 Introduction

Please accept my resignation. I don't want to belong to any club that will accept people as me as their member.

Groucho Marx, message to the Friar's Club of Beverly Hills

The preceding quote demonstrates a "winner's curse". Obtaining a membership reveals to Groucho Marx that it actually is undesirable: winning is bad news. It also reveals a failure of contingent thinking. When the comedian realizes the club's udesirability, he resigns. But had he thought through all possible contingencies beforehand he should realize that applying only makes a difference in the scenario that played out. A natural question then is: why did he apply in the first place? In other words, what causes the failure of contingent thinking in this winner's curse environment. This paper attempts to provide an answer.

The possibility for a *winner's curse* (WC) is present in a broad range of economic situations, including common-value auctions, financial markets and bargaining. Since Capen et al. [1971] first discussed the WC, economists have accumulated evidence supporting its existence, primarily in auctions[1] but also in other environments[2][3]. Charness and Levin [2009] are perhaps the first to design an experimental study that tries to uncover the drivers of the WC. Using different variations of the Acquiring a Company game ([Samuelson and Bazerman, 1984]) they find that the WC reflects individuals diffulties with performing contiget thinking. Following their work, a substantial body of work has formed documenting the difficulties people face to behave optimally in situations that require contingent thinking.[4] Notably, these go far beyond common-value games of incomplete information, like the auctions where the WC was first observed. A *failure of contingent thinking* (FCT) can also manifest in games with complete information, as well as in non-strategic individual choice problems. Niederle and Vespa [2023] discuss these issues, along as several different models that have been used to explain these FCTs.

Given the origins of the WC concept and its relationship to FCT, it is somewhat surprising that little work has been done on studying FCT in common-value environments

[1]See Mead et al. [1983], Hendricks et al. [1987] for evidence from the field and Kagel and Levin [1986], Kagel et al. [1987, 1995], Kagel and Levin [2002] for experimental auctions.

[2]See for instance Bazerman and Samuelson [1983], Cassing and Douglas [1980], Blecherman and Camerer [1996], Roll [1986], Rock [1986], Levis [1990].

[3]Thaler [1988] and Kagel and Levin [2002] provide comprehensive surveys

[4]See Shafir and Tversky [1992], Friedman [1998], Rabin and Weizsäcker [2009], Louis [2012], Esponda and Vespa [2014, 2018], Cason and Plott [2014], Eyster and Weizsacker [2010], Enke and Zimmermann [2019], Araujo et al. [2021], Moser [2019], Ngangoué and Weizsäcker [2021], Park [2025].

with incomplete information.[5] As noted by Niederle and Vespa [2023], it remains an open question whether an FCT in these situations is captured by the other known cognitive biases like a failure to correctly update beliefs (e.g. [Eyster and Rabin, 2005]) or it is an independent phaenomenon, and indeed one that has not yet been appropriately formalized. Here I attempt to shed some light on this issue.

I use what I term the *Barrel of Apples game* (BoA). This is a common-value game of incoplete information where one player is exposed to a potential winner's curse.[6] The game can be played either simultaneously or sequentially and the payoff-maximizing choice for the player exposed to the WC is always the same. Intuitively, one expects fewer mistakes to be made in the sequential version of the game where contingent thinking is not necessary. In an experiment reported in Louis [2012], participants played both the simultaneous and sequential versions of the BoA. In both cases the majority failed to detect the payoff maximizing strategy. Nevertheless, substantially fewer participants suffered from the WC in the sequential version of the game compared to the simultaneous one.[7]

There are different mechanisms related to cognitive computations that could have this effect on behavior. Eyster and Rabin [2005] propose a theoretical model that explains *cursedness* – i.e. suffering from the WC – purely as a failure of individuals in making inferences from others' actions.[8] In the conclusions they state (emphasis added):

> *"Another line of generalization would be to add more realistic variation in the degree of cursedness in different situations. For instance,* ***players are probably more likely to ignore the informational content of other players' actions***

[5]Ivanov et al. [2010] design a clever experiment that uses a second-price common-value auction environment and find no strong support for belief-based explanations for the WC. However, they do not test for FCT. The experiments in Li [2017] compare behavior under different mechanisms where the environment is either private-value with incomplete information, or common-value with complete information. Moser [2019] is perhaps the only study of the ralationship between FCT and the WC using a common value auction setup.

[6]The game is not an auction. Nevertheless it can be viewed as a very simple binary allocation problem. An example of a related task in a common-value environment with complete information are the random serial dictatorship games in Li [2017].

[7]See also Esponda and Vespa [2014] that find very similar results in a variation of the voting game they use. For more general comparisons of behavior in simultaneous and sequential mechanisms see Kagel and Levin [2002] and Li [2017] for auction environments and Ali et al. [2008] for committee voting.

[8]Analogy based expectation equilibrium (Jehiel 2005 Jehiel [2005], Jehiel and Koessler 2008 Jehiel and Koessler [2008]) is a closely related but more general concept. Correlation neglect ([Eyster and Weizsacker, 2010, Enke and Zimmermann, 2019] would also work in a very similar way as cursedness in situations where the other player is nature. Cognitive Hierarchy and Level-k type models ([Stahl and Wilson, 1994, 1995, Nagel, 1995, Camerer et al., 2004]) are another possibility. Such step-level reasoning can explain the winner's curse in some contexts ([Gneezy, 2005, Crawford and Iriberri, 2007], but not the change in behavior between simultaneous and sequential games.

when they have not actually observed these actions than when they have; *observing actions seems likely to induce more strategic sophistication. Hence,* ***players in certain sequential games may be less cursed than they would be in corresponding simultaneous-move games.***"

Thus, according to this interpretation the FCT observed in the BoA game is simply a result of cursedness, namely the failure to associate others' actions to their type. In other words, the ability to make inferences from others' actions is dependent on whether these actions are actually observed.[9] If this is the case, then players should not react to any changes in the private information structure of others. For instance, observing that another bidder acquires additional information before a common-value sealed-bid auction should not affect one's bidding strategy.

Still, there is a different possibility. Players may always be conscious of the correlation between others' actions and their private information — not cursed, in Eyster and Rabin's terms — but unable to perform correct contingent thinking.[10] That is, they understand that observing others' actions can reveal something about their private information, but fail to correctly represent all relevant hypothetical states in their mind, leading to sub-optimal behavior. Again, if this is the case then any change in an opponent's private information structure should have an effect on one's own decisions.

The main contribution of the paper is to distinguish between these two potential mechanisms for explaining the WC. To do so I use an experiment where subjects play the Barrel of Apples Game simultaneously, sequentially and in an extended version dubbed 'pay to observe'. I employ a within-subject design and classify subjects into behavioral types using their behavior in the simultaneous and sequential games. By correlating behavior in the 'pay to observe' version with subjects' behavioral types we can draw conclusions about what cognitive computations appear to drive behavior. This is possible by manipulating one player's information structure and observing the effects on the other's decisions.

I find that about a third of the subjects, like Groucho Marx, perform better in the sequential compared to the simultaneous game. The remaining subjects' behavior seems unaffected by the game's format and are split equally between naïve and sophisticated types. More importantly, behavior in the extended game indicates that Groucho Marx

[9]See Cohen and Li [2022] and Fong et al. [2023] for some attemts to formalize this notion. More generally, the possibility of individuals' level of sophistication being dependent on features of the environment has been studies both theoretically (e.g. [Li, 2017, Alaoui and Penta, 2022]) and experimentally (e.g. [Agranov et al., 2012, Alaoui and Penta, 2016]).

[10]See Piermont and Zuazo-Garin [2020] who provide a general theoretical framework for analyzing this type of failures.

types are always able to make inferences from others' actions, even when these are hypothetical. The change in behavior between the two game formats should be attributed to a failure of contingent thinking by those types. Only naïve subjects seem to be truly unable to make inference from others' actions and changes in behavior cannot be explained by the mechanism that Eyster & Rabin conjecture.

To test the robustness of my findings I also use an alternative approach. I fit to the data structural models of reasoning that allow for mixtures of reasoning modes. In the best fitting model, the estimated shares of types differ slighlty. The majority of subjects is estimated to be cursed, while between a quarter and a third of subjects are estimated to suffer from a pure failure of contingent thinking. Hence, myy main findings appear to be robust.

The paper proceeds as follows: I first introduce the Barrel of Apples game and discuss the theoretically expected behavior, as well as some reasoning models that might explain it. I then describe the experimental design and the associated research plan and hypotheses. The next section provides the experimental results and analysis thereoff. In the final section I provide some discussion and conclusions.

2 Theoretical framework

My prime aim is to uncover the cognitive mechanism that gives rise to the winner's curse in an unabiguous way. It is therefore best to pursue this in a decision environment that features a winner's curse and is very simple, yet flexible enough to allow us to disentangle the role of different mental computations experimentally, even if said environment does not directly match any known economic setup. I therefore use the *barrel of apples game*, first introduced in Louis (2012).

2.1 The barrel of apples game

In the barrel of apples game two players indexed by $i \in \{1, 2\}$ are offered a barrel of 10 apples. The barrel may be *good*, in which case it contains 10 good apples, or bad, which means it contains a number $Q \in \{1, \ldots, 9\}$ of *bad* apples. I denote the state, *good* or *bad*, as $\theta \in \{G, B\}$. Both states are equally likely. Both players draw an apple from the barrel, check its quality and put it back. Notice that drawing a good apple means it is more likely for the barrel to be good. Drawing a bad apple means the barrel is certainly bad. Formally, the apple drawn constitutes each player's private signal about the state denoted by $s_i \in \{g, b\}$. Then $Pr\{\theta = 1 \,|\, s_i = g, Q\} > \frac{1}{2}$ and $Pr\{\theta = 1 \,|\, s_i = b, Q\} = 0$, for any Q.

Both players choose whether to accept or reject the barrel: $x_i \in X = \{\mathbf{A}ccept, \mathbf{R}eject\}$. It is not possible for both players to share the barrel. To resolve conflict in case both choose to accept, player 1 is assigned priority: if both players accept the barrel, only player 1 obtains it. In other words, player 2 can obtain the barrel only if player 1 rejects it. It is possible for no player to obtain the barrel. This is the case when both players reject it. Formally, let $f_i : X^2 \to X$ be the assignment function for player i. Then we have[11],

$$f_1(x_1, x_2) = x_1$$

and

$$f_2(x_2, x_1) = \begin{cases} x_2, & for\ x_1 = R \\ R, & for\ x_1 = A \end{cases}$$

where, for notational economy I use the set of choices to also denote the set of outcomes: $f_i(x) = A$ means player i gets the barrel and $f_i(x) = R$ means he does not get it.

The players' payoffs depend on the state and whether they obtain the barrel, but not on Q. Players want to obtain a good barrel and avoid a bad barrel. A player gets a payoff V if she *obtains a good barrel* or *does not obtain a bad barrel*. In all other cases the payoff is 0.[12] Let the function $u_i : \{A, R\} \times \{G, B\} \to \mathbb{R}$ represent player i's payoff. We have:

$$u_i\left(f_i(x_i, x_{-i}), \theta\right) = \begin{cases} V, & f_i(x_i, x_{-i}) = A \ \&\ \theta = G \\ V, & f_i(x_i, x_{-i}) = R \ \&\ \theta = B \\ 0, & otherwise \end{cases}$$

I consider three variations of the game: two with a single stage and one with two stages. In the *simulatenous* format both players make decisions simultaneously. In the *sequential* format player 1 makes a decision first, followed by player 2. Player 2 observes player 1's choice, but not her private signal when making his choice.[13] Finally, in the *pay to observe (PtO)* format there are two stages. In the first, player 2 has the option to pay a small price to be able to observe player 1's choice. The decision to pay to observe is made before player 2 can observe his own signal. Having the option to pay to observe player 1's choice essentially means that player 2 can choose whether to play

[11] This formalization is useful for the comparison of the different reasoning models presented further on

[12] Notice that by setting up payoffs in this way there is no risk safe action for neither player. This is done to control for risk aversion when the game is used in an experiment.

[13] For ease of exposition I will use female pronouns for player 1 and male pronouns for player 2.

the simultaneous format of the game in the second stage or pay a small price to play the sequential format.

The game falls within the class of binary allocation problems, although it is clearly not an auction. Perhaps it can be described as a serial dictatorship mechanism.

It should be clear that the BoA game is quite simple. Both players are faced with a binary choice. What's more, Player 1's outcome does not depend in any way on the choices of Player 2. In fact, as we see next, Player 1 has a very simple and intuitive payoff-maximizing choice strategy. This should simplify the formation of correct beliefs for Player 2, although as we will see this is likely not the case. Another advantage this has is that Player 1 can easily be substituted by a computer endowed with this simple and intuitive choice rule. This is something I leverage in this experiment, although in Louis [2012] both players were human. All humans that played in the role of Player 1 there behaved optimally.

The *simultaneous* and *sequential* formats are used in the experiments of Louis [2012], while the *PtO* variation is novel. Following the results there, the research strategy here is to first use behavior in the former two formats to classify subjects into behavioral types, and then to compare the choices of these types in the *PtO* format.

2.2 Expected behavior

Given player 1's priority over player 2, her payoff depends only on the state of nature and her own actions. This is true in all three formats of the game that I consider. Hence, the best strategy is for her to always follow her private signal: if the apple drawn from the barrel is bad, reject the barrel; if the apple is good, accept. For the remainder of the paper I assume that player 1 follows this optimal behavior rule and focus attention on the behavior of player 2.[14]

First, let us consider the case where the two players make their choices *simultaneously* assuming that player 1 plays optimally, player 2 faces a *winner's curse*: he can only obtain the barrel when player 1 rejects; player 1 rejects only after observing a bad signal; a bad signal reveals perfectly that the barrel is bad. It is therefore optimal for player 2 to always reject the barrel, irrespectively of his own private signal.

The above reasoning is still valid if the game is played *sequentially* with player 1 playing first. The difference here is that player 2 gets to make a choice only when player 1 has rejected the barrel, and therefore revealed through her choice that her private

[14] As I mention above and explain in detail later, player 1 in the experiment is played by a computer that follows exactly this optimal behavior rule.

signal and the barrel are bad. Still, it is again optimal for player 2 to always reject the barrel, irrespectively of his own private signal.

In the *PtO* format player 2 decides whether the game is played simultaneously or sequentially. Once that decision is made, there is no difference from the other two formats and the discussion above still applies: it is optimal for player 2 to always reject the barrel. Since the optimal choice is not affected by the game's format, there is no value for player 2 in observing player 1's choice. Hence it is optimal to never pay to observe. Note that in *PtO* player 2 chooses whether to pay to observe before receiving his private signal, but this would not make a difference here.

The discussion above is assuming that an individual in the role of player 2 possesses some level of sophistication that allows it to think through the situation in the way described. In the next section I discuss particular types of reasoning that can give rise to behavior that deviates from the sophisticated benchamrk described. Before doing so it is useful to classify expected behavior in the different game formats.

One type of non-optimal behavior I expect to find is that of individuals that decide whether to accept or reject the barrel based on their private signal in both the simultaneous and sequential versions of the game. I dub such behavior as *naive*. An individual behaving like that is not influenced by the format of the game and therefore is not expected to see any value in playing the game sequentially instead of simultaneously. Hence individuals behaving naively in the *simulatnesous* and *sequential* formats are expected to not pay to observe in the *PtO* format.

While sophisticated behavior posits that an individual should reject the barrel in both single stage formats of the games, reaching this conclusion in the *simultaneous* format requires reasoning about hypothetical events. I therefore expect some individuals to behave optimally in the *sequential* format, but to follow their signal in the *simultaneous* one. Following Louis [2012] I call this GM behavior which stands for Groucho Marx.[15] If an individual's behavior is influenced by the game's format, it is possible that it also finds some value in observing player 1's choice. I therefore expect dispalying GM type behavior to also pay to observe in the *PtO* format.

The expected types of behavior described here are summarized in Table 1. The experiment is designed in a way that allows us to distinguish whether subjects's behavior conforms to these predictions. It also allows us to obtain some insight into the cognitive biases that give rise to such behavior. These are discussed next.

[15] See the discussion in the Introduction

Behavior	1. Simultaneous	2. Sequential	3. Pay to observe
Sophisticated	Reject	Reject	No
Naive	Follow signal	Follow signal	No
GM	Follow signal	Reject	Yes

Table 1: Expected types of behavior

2.3 Models of cognitive reasoning

2.3.1 Rational expectations

It is straightforward to see that in the absence of an unusual form of preferences, forming rational expectations should lead to the sophisticated type of behavior described in the previous section. In the context of the barrel of apples game, rational expectations imply that player 2 forms beliefs about player 1's choice and incorporates them correctly into the calculation of expected payoffs. Let $\sigma(x_1|s_1) \in [0,1]$ represent player 2's beliefs about the probability of player 1 choosing x_1 after receiving a signal s_1. It must be:

$$\sigma(R|s_1) = \begin{cases} 1, & s_1 = b \\ 0, & s_1 = g \end{cases}$$

Then player 2 's expected payoff from choosing x_2 after observing a signal s_2 in the simultaneous format is given by:

$$U_R^{sim}(x_2|s_2) = \sum_{\theta \in \{G,B\}} Pr\{\theta|s_2\} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \sum_{x_1 \in \{A,R\}} \sigma(x_1|s_1) \cdot u_2\left(f_2(x_2,x_1),\theta\right) \quad (1)$$

When the game is played sequentially player 2 only plays when $x_1 = R$, which he observes. Expected payoffs in this format are calculated as follows:

$$U_R^{seq}(x_2|x_1 = R, s_2) = \sum_{\theta \in \{G,B\}} Pr\{\theta|x_1 = R, s_2\} \cdot u_2\left(f_2(x_2,x_1),\theta\right)$$

After applying the Bayes theorem and rearranging the terms, the above can be written as:

$$U_R^{seq}(x_2|x_1 = R, s_2) = \sum_{\theta \in \{G,B\}} \frac{Pr\{\theta\}Pr\{s_2|\theta\}}{Pr\{x_1 = R, s_2\}} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \cdot \sigma(x_1|s_1) \cdot u_2\left(f_2(x_2,R),\theta\right) \quad (2)$$

It is straightforward to use expressions (1) and (2) to conclude that it is optimal for a rational individual to always reject the barrel. Furthermore, a rational individual understands that there is no value in oberving player 1's choice and therefore never pays to observe her signal in the PtO format of the game.

But the purpose writing down player 2's expected payoffs explicitly is different. If we want to understand what causes behavior to deviate from what is optimal, it makes sense to look at models of reasoning that may give rise to non-rational expectations. The above expressions allow us to highlight how such alternatives differ from the benchmark and from each other.

2.3.2 Cursedness or correlation neglect

A first such candidate can be found in the cursedness bias undelying the cursed equilibrium concept ([Eyster and Rabin, 2005]). The main idea is that individuals may fail to take into account how others' actions depend on their information. In this specific context I assume that when a cursed player 2 calculates his expected payoffs, he fails to incorporate into these calculations the correlation between player 1's choices and her private signal.[16] Instead, his beliefs about the probability of player 1 accepting the barrel equal the average probability of this happening and are independent from the players' private signals.[17] Formally, let $q = \frac{Q}{10}$ and the average beliefs be $\overline{\sigma}(x_1 = A) = Pr(x_1 = A) = \frac{2-q}{2}$. Cursed beliefs are given by $\sigma^C(x_1|s_1) = \chi \cdot \overline{\sigma}(x_1) + (1-\chi) \cdot \sigma(x_1|s_1)$. In principle, the parameter χ can take any value in the unit interval. For the purposes of this paper I focus on the extreme values. For now let us consider a fully cursed individual with $\chi = 1$, so $\sigma^C(x_1|s_1) = \overline{\sigma}(x_1)$. The expected payoffs for such a fully cursed player 2 in the simultaneous format are given by:

$$U_C^{sim}(x_2|s_2) = \sum_{\theta \in \{G,B\}} Pr\{\theta|s_2\} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \sum_{x_1 \in \{A,R\}} \sigma^C(x_1|s_1) \cdot u_2\left(f_2(x_2, x_1), \theta\right) \tag{3}$$

[16]Notice that I do not imply that a cursed individual is necessarily unable to predict player 1's choice. The problem lies in using such ability appropriately when forming expectations about one's own payoffs, that are contingent on others' choices. This is an important distinction given the experimental design and I discuss it further after presenting the design in the next section.

[17]Eyster and Rabin's (2005) specification of cursedness differs in two ways. First, average beliefs are allowed to depend on a player's own type. Second, the beliefs used by a player are a convex combination between average beliefs and a correct prediction of others' play. I only consider fully cursed players. These deviations from the original model maintain the main features of cursedness, simplify exposition, and do not affect my main results.

And in the sequential format by:

$$U_C^{seq}(x_2|x_1 = R, s_2) = \sum_{\theta \in \{G,B\}} \frac{Pr\{\theta\}Pr\{s_2|\theta\}}{Pr\{x_1 = R, s_2\}} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \cdot \sigma^C(x_1|s_1) \cdot u_2\left(f_2(x_2, R), \theta\right) \quad (4)$$

Such a fully cursed player 2 finds it optimal to follow his private signal in both the simultaneous and sequential formats, giving rise to a naive behavior pattern. Also, as such a player 2 cannot incorporate the correlation of player 1's decision with her signal, he finds no value in observing her choice, and neither for paying to do so in the PtO format of the game.

The concept of cursedness is based on the idea that one faces a human opponent. One could extend this model to situations where the opponent is nature, i.e. the individual is unable to make inferences about the state of the world from observing nature's moves. This can be then viewed as a version of correlation neglect ([Eyster and Weizsacker, 2010, Enke and Zimmermann, 2019] and would give the exact same results. Given that in the experiment Player 2 faces a computer it might be argued that referring to this as a model of correlation neglect is more appropriate, as there are no beliefs about another human player involved. Whether individuals form beliefs about humans differently than they do for machines is an empirical question that has acquired increased relevance given recent technological developments, but not one I can address here.[18] I will use cursedness throughout, meaning a failure to take into account the correlation between states of the world and actions, regardless of whether these actions are taken by a human, a machine, or nature.

2.3.3 Varying Cursedness

As an alternative to the simple specification of cursedness I just described, I also consider the possibility of the level of cursedness, captured by χ, to vary depending on the game's format. In particular I look at the case where $\chi = 1$ in the simultaneous format and $\chi = 0$ in the sequential one. This is in line with the generalization proposed in Eyster and Rabin (2005) (see quote in the Introduction). It is also inline with attempts to extend the cursed equilibrium model to sequential games ([Cohen and Li, 2022, Fong et al., 2023]). In this case we have:

$$U_{VC}^{sim}(x_2, s_2) = U_C^{sim}(x_2, s_2) \quad (5)$$

[18] I should note that in the experiment in Louis [2012], where humans faced humans, behavior was very much in line with what I find here. Of course there were other important differences in the design of the two experiments as well.

and

$$U_{VC}^{seq}(x_2|x_1 = R, s_2) = U_R^{seq}(x_2|x_1 = R, s_2) \tag{6}$$

Given the above calculations, a player 2 with varying cursedness is expected to exhibit a GM type behavior in the single-stage formats of the game. Nevertheless, the decision to pay to observe player 1's choice in the PtO format is made before observing anything. Hence, at this stage the individual remains fully cursed. It is not possible for him to anticipate the change in cursedness (and behavior) that should follow after observing player 1's choice. Thus, similar to a fully cursed player 2, one that diplays varying cursedness should also never pay to observe in the PtO format.

2.3.4 A pure Failure of Contingent Thinking

I consider one final possibility. That is of an individual that, on the one hand, is not cursed and can perfectly incorporate the correlation between player 1's signal and her choices when calculating expected payoffs. On the other hand, that individual has difficulties in thinking contingently about the outcomes. In particular, it cannot incorporate correctly the way the allocation depends on the other player's choice. Formally, I assume that instead of using the game's allocation function $f_2(\cdot)$, a non-contingent thinker (NCT) uses $f_2^{NCT} : X \to X$, which has the following form:[19]

$$f_2^{NCT}(x_2) = x_2$$

Expected payoffs are then as follows:

$$U_{NCT}^{sim}(x_2, s_2) = \sum_{\theta \in \{G,B\}} Pr\{\theta|s_2\} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \sum_{x_1 \in \{A,R\}} \sigma(x_1|s_1) \cdot u_2\left(f_2^{NCT}(x_2), \theta\right) \tag{7}$$

and

$$U_{NCT}^{seq}(x_2|x_1 = R, s_2) = \sum_{\theta \in \{G,B\}} \frac{Pr\{\theta\}Pr\{s_2|\theta\}}{Pr\{x_1 = R, s_2\}} \sum_{s_1 \in \{g,b\}} Pr\{s_1|\theta\} \cdot \sigma(x_1|s_1) \cdot u_2\left(f_2^{NCT}(x_2), \theta\right) \tag{8}$$

First, notice that since $f_2^{NCT}(x_2) = f_2(x_2, R)$, it follows from (2) and (8) that

$$U_{NCT}^{seq}(x_2|x_1 = R, s_2) = U_R^{seq}(x_2|x_1 = R, s_2)$$

[19]One can of course consider a more general model in which player 2 believes he will get the barrel with some probability $\phi(x_2, x_1)$, which is non-decreasing in x_2 and non-increasing in x_1. For the purpose of this paper this simple version is sufficient.

Therefore, the NCT player 2 correctly calculates $x_2 = R$ to be optimal in the sequential format. On the other hand, in the simultaneous format he erroneously finds that following his signal is optimal. Unlike the case for cursed players, here player 2 understands that player 1 only rejects after observing a bad signal. The mistake lies in calculating as if he can always get the barrel by playing $x_1 = A$, even when player 1 accepts it. He fails to see that the only contingency that should matter for calculating expected payoffs is the one in which player 1 rejects. Thus, an NCT player 2 displays GM type behavior in the single-stage versions of the game.

The crucial difference of the NCT model from all the others here is its prediction about behavior in the PtO format of the game. Notice that unlike the varying cursedness model, an NCT player 2 remains NCT regardless of the choice he makes in the first stage. Hence, he can very well anticipate his change in behavior depending on whether the game is played simultaneously or sequentially. In particular his "optimal" choices in the simultaneous game depend only on his signal:

$$x_2^{sim}(s_2) = \begin{cases} A, & s_2 = g \\ R, & s_2 = b \end{cases}$$

In the sequential game his choice depends on his signal and the choice of player 1:

$$x_2^{sim}(s_2, x_1) = \begin{cases} A, & s_2 = g \ \&\ x_1 = A \\ R, & s_2 = b \ or \ x_1 = R \end{cases}$$

More importantly, he understands that observing player 1 rejecting the barrel means that he should also reject it. He therefore finds there to be value in such an observation and might be willing to pay a price to make the game sequential. This value is given by (recall that player 2 does not observe his own signal when deciding whether or not to pay to observe):

$$U_{NCT}^{seq}\left(x_2^{seq}(s_2, x_1)\right) - U_{NCT}^{sim}\left(x_2^{sim}(s_2)\right) \quad = \quad \frac{V}{2}q(1-q)\ , \quad \left[where\ q = \frac{Q}{10}\right] \tag{9}$$

This value is always positive. Interestingly, in depends non-monotonically on Q, the number of bad apples in the bad barrel. In particular, the value is maximized for the intermediate value $Q = 5$. The intuition here is that an NCT player 2 understands that player 1's choice perfectly reveals her private signal. Therefore, the desicion to pay to observe her choice is essentially a choice of acquiring an additional draw (with replacement) from the barrel. The marginal value of the second draw depends on the signal's marginal informativeness. If $Q = 0$, there is no information, and therefore no value, in either of the private signals. If $Q = 10$ signals would be perfectly informative, so the additional signal can add no value. The marginal informativeness of the 2nd signal is higher for intermediate values of Q, which results in the non-monotonic relationship given in (9).

At a conceptual level, the notion of a player type that could not think contingently was the motivation for the experimental design used here. In particular, I hypothesized that such a type would exhibit GM type behavior in the single-stage formats of the game and be willing to pay to observe player 1's choice in PtO. Previous versions of this paper included non-formal descriptions of such a type and its predicted behavior. The formal specification presented here was developed after observing the experimental results. It incorporates in a simple way the hypothesized conceptual differences between the reasoning of such a type and a (varyingly) cursed player, as well as the differences in their predicted behavior. The particular specification has the added benefit of capturing the non-monotonicity between the value an NCT player assigns to observing player 1's choice and Q, for which I also find evidence in the experimental results, but did not anticipate when designing the experiment.

3 The experiment

3.1 Experimental Design

I first give a detailed presentation of the experimental design. Some of the design choices may warrant further discussion and justification. This is provided at the end of this section.

The experiment employs a fully randomized within-subject (or repeated measures) design. Subjects play variations of the barrel of apples game repeatedly, without feedback, for 40 rounds taking the role of Player 2. Player 1 is simulated by a computer endowed with a decision rule that mimics Player 1's optimal choice: i.e. accept when drawing a good apple, reject when drawing a bad apple. Subjects are informed about the computer's choice rule.

Subjects face all three game formats in the 40 rounds of play in random order. They play the simultaneous and PtO formats in 10 rounds each, and the sequential format in 20 rounds. The sequential format is played twice as many times to ensure a sufficient number of observations. Each time the computer accepts the barrel in this scenario, the game ends without a decision from the subjects. On average, subjects make about five decisions in the 20 sequential conditions.

The number Q (the number of bad apples in the bad barrel) varies between 1 and 9. In particular, it takes value 5 two times in the simultaneous and PtO formats, and four times in the sequential format. All other values are taken once in the former two and twice in the latter. The combination of a particular format with a specific number of Q results in a specific condition. Table 2 shows all possible conditions and the number of times each subject is in one of them. The order in which subjects go through the conditions is randomized separately for each subject. No feedback is provided to subjects after their choices in each condition until the end of the experiment.

	Q , # of bad apples									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**Total**
1. Simultaneous	1	1	1	1	2	1	1	1	1	10
2. Sequential	2	2	2	2	4	2	2	2	2	20
3. Pay to observe	1	1	1	1	2	1	1	1	1	10

Table 2: Number of decisions per condition for each subject

Apples are presented as balls in urns on the subjects' screens. They have different colors depending on whether they are good or bad. The two possible colors are blue and red, but their meaning (good/bad) is not fixed and changes randomly in each game.

The prize V for either obtaining the good barrel or not obtaining the bad barrel is set at 100 points. In all other cases, the payoff is zero. In the PtO, the price to observe Player 1's choice is 10 points. To determine subjects' earnings, 7 of the 40 rounds were chosen randomly and the subjects' outcome in these rounds was considered. An additional 100 points were given to each subject to ensure that there was no bankruptcy from the choices in the PtO format. Points were converted to Swiss francs using a 20 to 1 exchange rate. Each subject received an additional 10 francs show-up fee.

Detailed instructions were handed out to subjects at the beginning of the experiment and subjects were left with enough time to read through them and ask clarifying questions. The exact text of the instructions can be found in the appendix. After reading the instructions and before starting the experiment, subjects had to pass a short comprehension test.[20] They were not allowed to continue until answering all questions correctly. Subjects had no significant difficulties in doing so.

[20]The five questions can be found in the Appendix.

The experiment was conducted in six sessions with a total of 97 subjects participating. This gives 97 independent clusters of observations, as there is no interaction among participants. While the theoretical maximum number of observations per participant is 40, the actual number is lower, since whenever the game is played sequentially, i.e. in *Seq* or after paying in *PtO*, and the computer accepts the barrel, the participant makes no choice. In total we observe 2,155 *Accept/Reject* choices and 970 *pay to observe* choices.[21] Subjects were recruited among undergraduate students at ETH Zurich and the University of Zurich using ORSEE (Greiner 2003 Greiner [2004]). The experiment was conducted in the Experimental Economics Lab of the University of Zurich using Z-Tree (Fischbacher 2007 Fischbacher [2007]).

I now provide some justification for the design choices described above, starting from the choice of a within-subject design. My main research question requires me to classify participants into different types. This requires observing a participant's behavior in different formats of the BoA game. The variation of Q is necessary to understand if and why different cognitive types behave differently in the *PtO*, which in turn may inform us about the origin of the failure in contingent thinking. For these reasons a within-subject design was necessary. An added benefit from this choice is that such a design offers increased power to detect true treatment effects compared to between-subject designs. The reason is that any systematic individual differences do not contribute to the error term (Vonesh, 1983 Vonesh [1983], Maxwell and Delaney, 2003 Maxwell and Delaney [2003]). An important consideration when using within-subject designs are possible order effects. Fully randomizing conditions separately for each subject provides a means of controlling for such effects in a probabilistic manner (see Maxwell and Delaney, 2003 Maxwell and Delaney [2003]).

Not providing feedback after each round is another design choice dictated by the research question at hand. Classifying participants into types relies on the assumption that behavior in the different formats is relatively stable across rounds. Feedback about outcomes from earlier rounds would most likely affect behavior in later rounds, making classification very hard. That is not to say that feedback may not play a role in mitigating the FCT, just that the research question considered here is better served by controlling for this possible confound.[22]

[21] The number of participants per session was: 5, 15, 15, 20, 21, 22. This number varied because of different show-up rates. In particular, the low number of participants in the first session is attributable to a failure of the recruitment system to send the automated reminder to subjects. A pilot session conducted during the calibration of the design is not included in the analysis. The results I present are robust to omitting any particular session.

[22] The existing studies on the FCT show that indeed feedback can help some individuals learn to play

Besides the reasons described above, the random variation of conditions across rounds, including the "meaning" of the balls' color, paired with the lack of feedback, has the added benefit of forcing participants to view every round as a new decision problem. In other words, it should be clear to them that they cannot rely on simple heuristics (e.g. "always reject when the ball is red"), but have to employ cognitive effort and rely on their underlying reasoning abilities for each one of their choices. As long as the incentives are sufficiently salient, this should reveal the differences in their ability to perform cognitive thinking, which is crucial for this research design.

Designs with such a high degree of variation across rounds are not common in economic experiments and one might be worried that the constant variation across rounds might lead to confusion.[23] Of course, the clear instructions and comprehension test at the start are employed in an effort to mitigate such a possibility but here is no guarantee that they succeed. Ultimately, it is an empirical question whether the behavior we observe in the experiment is driven by confusion. For example, in this design, seeing a bad signal perfectly reveals the state being bad. One might expect confused subjects to sometimes accept even in this case. As I show in the next section, this happens very rarely in the experiment. More generally, as I show next, the choices subjects make in the *Sim* and *Seq* formats are in line with what other studies in the literature find.

3.2 Research plan and hypotheses

The aim of the experiment is to provide insight into the mechanism leading to a failure of contingent thinking. Existing experimental work, (e.g. Louis, 2012 Louis [2012] or Esponda and Vespa, 2014 Esponda and Vespa [2014]) demonstrates both the difficulty individuals have to make optimal choices when contingent thinking is required, as well as the impovement that can be achieved by placing these individuals in the only relevant contingency. The goal here is not to simply replicate these existing results. Nevertheless, being able to do so is a prerequisite for the remaining steps of the research plan.

Hypothesis 1: I expect to observe on average more optimal behavior in choices made under the *sequential* format compared to the *simultaneous* one.

---

optimally, but the FCT is far from eliminated through learning with feedback. See Niederle and Vespa (2023) Niederle and Vespa [2023] for a more detailed discussion of the topic.

[23]Fully randomized within-subject designs are much more common in psychology and neuroscience experiments. Perhaps more pertinent to note is the fact that the experiments in Esponda and Vespa (2014) Esponda and Vespa [2014] employ comparable levels of variation in parameters across rounds for similar reasons.

Beyond these aggregate level obervations, I follow a two step research plan. In the first step I use behavior in the *Sim* and *Seq* conditions to classify them into behavioral types. Based on previous results in Louis (2012) Louis [2012] I expect the vast majority of participants to follow of the three behavioral patterns described previously whenever they make a choice to accept or reject the barrel (see the first three columns of Table 1).

Hypothesis 2: When choosing to accept or reject the barrel, individual behavior will fall under one of the three types: *Sophisticate*, *Naive* and *GM*.

In the second step of the research plan I focus on the behavior in the first stage of the *PtO* conditions, namely the choice of whether or not to pay to observe Player 1's choice. This is the most important part of the research plan and the one that can provide novel isights into the relationship between the FCT and cursedness (or correlation neglect). I compare the the behavior across the different types obtained in the previous step. I expect this to be in line with what is shown in the last column of Table 1.

Hypothesis 3: When choosing whether to pay to observe the computer's choice, only *GM* type individuals will do so. Both *Sophisticate* and *Naive* type individuals will not pay to observe the computer's choice.

To test the robustness of my findings I leverage the models of reasoning introduced earlier. I use the data to estimate a series of different structural models that allow for a mixture of different modes of reasoning. This fitting exercise uses the complete set of data, namely both the Accept/Reject decisions from all three formats, as well as the pay to observe choices made in the *PtO* conditions. Comparing the different models allows us to infer what types of reasoning seem to best describe how participants in the experiment think.

4 Results

The results of the experiment are presented as follows. I first provide some agregate results on behavior across the different formats. Then I follow the steps of the research plan laid out by first: classifying subjects into behavioral types using data from the *Sim* and *Seq* conditions; second: comparing behavior in *PtO* across types.

4.1 Aggregate results

Table 3 shows the aggregate behavior of subjects across treatments. The frequency with which subjects reject the barrel is indicated, conditional on being in each format, and after observing a good or bad signal. For *PtO*, these numbers are further broken down depending on whether the subject chose to pay to observe the computer's decision before making a choice to accept or reject the barrel.

	1. Simultaneous		**2. Sequential**		**3. Pay to observe**			
paid					*No* **62.3%**		*Yes* **37.7%**	
signal	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*
frequency of rejection %	**96.2**	**31.4**	**99**	**60.3**	**99.4**	**38.6**	**96.7**	**65.7**
# of obs.	235	735	297	189	153	451	60	35

Table 3: The frequency of rejection observed in the data, for the three game formats.

Recall that optimal behavior prescribes a 100% frequency of rejection, regardless of the format or one's private signal. On the one hand, participants almost always reject after observing a bad signal. This indicates them understanding that a bad signal unambiguously reveals the state to be bad. I interpret this as a sign that the continuous variation of conditions across rounds did not confuse the subjects substantially. On the other hand, the frequency of rejection does not take an extreme value in any of the conditions after observing a good signal. For the case of simultaneous decisions (in the second and sixth columns of Table 3) subjects reject the barrel after a good signal in about a third of the cases. For the case of sequential decisions (fourth and last column) the frequency of rejecting after a good signal roughly doubles.

The sensitivity of deviations from optimal behavior to the game's format and subjects' private signal is confirmed using probit regressions where I cluster standard errors at the individual level.[24] The regressions control for the number of 'bad balls' Q as well as the round in which a decision is taken. These are reported in Table 4. I normalize both Q and the round number to take values in the unit interval (denoted by q and *time* respectively) to keep coefficient estimates comparable.

[24]Panel data methods yield very similar results.

Probit regressions
Dependent variable: *subject's choice ('reject' = 0, 'accept' = 1)*
robust standard errors, clustered at the subject level

	Single-stage	PtO
subject's signal	2.325**	2.560**
	(0.180)	(0.247)
sequential	-0.758**	-0.579*
	(0.136)	(0.280)
q	0.610**	0.823**
	(0.194)	(0.254)
time	-0.343*	-0.501*
	(0.157)	(0.231)
intercept	-1.943**	-2.363**
	(0.233)	(0.293)
Log Lik	-630.03	-332.30
# of obs.	1456	699

* p-value < .05   ** p-value < .01

Table 4: Subjects' decision to accept or reject the barrel. Explanatory variables are: *subject's signal* (value 1 if the subject received a good signal, 0 otherwise); *sequential* (value 1 if the decision was made after observing the computer's choice. note: to record a decision by the subject in this case means the computer must have rejected); *q* (the fraction of bad apples in the bad barrel, takes values from .1 to .9); *time* (the normalized time in which a given decision is made, equals the round divided by 40, the total number of rounds, and takes values from 1/40 to 1).

In the first column of Table 4 I pool together obervations from the two single-stage formats. The second column uses the observations from the second stage of the PtO format. Notice that in the latter, the variable *sequential* is not exogenous, as it depends on the subject's decision in the first stage. It is interesting to see that there are small differences between the regression results in the two subsamples. I later argue that these are explained through the selection happening in the first stage.

The estimated coefficients for the *subject's signal* and *sequential* reflect what is shown in Table 3. The coefficient of q is positive and statistically significant. This is one reason that motivates the structural approach I use for classifying behavior next.

Finally, despite not receiving any feedback, it appears that subjects' behavior does get somewhat closer to optimal in later rounds. This suggests that subjects might be learning simply by repeatedly reasoning about the game, even without knowing the outcome of their reasoning.[25] One factor that might help such learning is the fact that subjects are often placed in the sequential format, which forces them to think through the only contingency that matters. This learning effect is not large, but it is statistically significant. I therefore take it into account when estimating structural models of behavior in section 4.4.

Overall, the results are in line with what previous studies find and support *Hypothesis 1*. Namely, the majority of participants find it difficult to choose optimally when contingent thinking is required. At the same time, when placed into the contingency where their choice matters for their payoff, the majority behaves optimally.

Result 1 The data supports Hypothesis 1. We observe on average substantially fewer optimal choices under the *Sim* format compared to *Seq*.

Before moving ahead with the main analysis of interest, which is that of individual behavior, it is worth briefly looking at the aggregate behavior in the first stage of the *PtO* conditions. This is the main novelty introduced in the current design. Recall that optimal behavior requires one to never "pay to observe". The first row of Table 5 shows the percentage frequency of participants paying to observe the computer's choice. Overall, we find that in slightly more than a third of cases subjects choose to pay to observe player 1's choice (see the right most cell of Table 5).

Turning attention to behavior in the first stage of the PtO format we find that in slightly more than a third of cases subjects choose to pay to observe player 1's choice (see upper right side of Table 3). Recall that the optimal choice here is to never pay. While this is the case in the majority of cases, overall behavior clearly deviates from this benchmark. An important observation here is that the choice in this case seems to be sensitive to Q. In fact, subjects choose to pay to observe player 1's choice more often when Q takes intermediate values, and less so when it takes values close to the extremes.

4.2 Classification

Each subject made 10 decisions in *Sim* with an average of 7.6 of those being made after receiving a good signal. Furthermore, each subject made an average of 5 decisions in

[25]See Weber 2004 Weber [2003].

Frequency of paying to observe (%)										
Q	1	2	3	4	5	6	7	8	9	Overall
All	**21.7**	**30.9**	**41.2**	**36.1**	**45.9**	**42.3**	**42.3**	**34.0**	**37.1**	**37.7**

Table 5: Decisions in the first stage of the PtO format. The numbers show the percentage of cases where a subject chose to pay to observe player 1's choice in the PtO treatments. The three last rows show the frequency broken down according to the classification described in section 4.2

Seq, with an average of 2 of those being made after receiving a good signal. For each individual we can calculate the frequency of rejection in each format, conditional on receiving a good signal.

The histograms in panels A and C of Figure 1 show that behaviour is heterogeneous in each of the two formats. In *Sim* the distribution is skewed to the left with a large proportion following their private signal and only a few behaving optimally. In *Seq* the distribution is bimodal with one mode in each extreme. What is more interesting though is the heterogeneity of a "higher order" which results from the product of these two distributions. This is depicted by the scatter plot in panel B.

As explained, our aim here is to classify participants into the behavioral types summarized in Table 1. The naive, sophistcated and GM benchmarks correspond to points (0,0), (1,1), and (0,1) respectively. A first observation is that while some subjects' behaviour fully conforms to these benchmarks, others are more noisy. This is not surprising given the high number of decisions made in the experiment. Other variables, such as Q, could explain some of the variation. A second, reassuring observation concerns the fact that very few observations lie in the lower right quadrant, an area that corresponds to the counterintuitive behaviour of rejecting in the simultaneous game and following one's signal after observing the computer rejecting. In fact, a significant mass of observations lies close to the points corresponding to the benchmark types and the spikes in the distributions directly coincide with these points. .

Result 2 The data largely supports *Hypothesis 2*. In *Sim* and *Seq* conditions a substantial mass of individuals behave in a way corresponding to one of the three behavioral types: *Sophisticated*, *Naive* and *GM*.

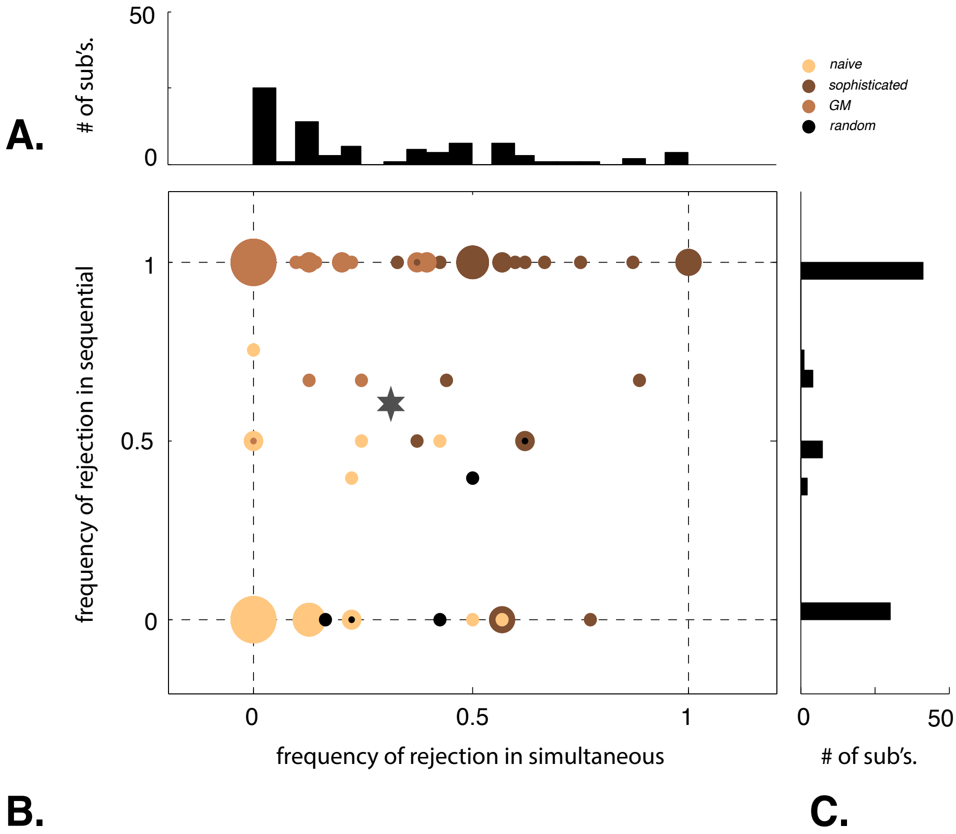

Figure 1: Individual frequency of rejection upon receiving a good signal, in *Sim* and *Seq* conditions. Each circle is centred at a point with coordinates corresponding to the frequency of rejection in the two formats. The radius of the circle is proportional to the number of subjects whose frequencies coincide. Colors refer to the types assigned by the classification scheme presented in subsection 4.3. The star indicates the average frequencies across all subjects. The points (0,0), (1,1) and (0,1) correspond to the naive, sophisticated and GM benchmarks respectively. The histograms above and to the right give the distribution of subjects over frequencies in *Sim* and *Seq* respectively.

Given the non-negligible "noise" in the data and to keep our our classification exercise robust, I allow for an additional *Random* type.

One way to classify subjects would be to use the euclidean (or some other measure of) distance of each data point in the scatter plot in figure 1 from the points corresponding to the benchmarks. While simple, this method is problematic for several reasons. The points in the scatter plot aggregate frequencies across treatments that differ with respect to an important variable, namely Q. This information would be lost if using distance only. Furthermore, there is an important number of observations with a frequency of rejection of exactly 50% in one or the other format. Classifying these subjects would require arbitrary choices on our part. Finally, 12 subjects did not make any sequential decision after receiving a good signal.[26]. It would therefore not be possible to classify these individuals using this method.

The method I use instead is distribution-based clustering. Choice distributions are derived from a hierarchy of logit choice models.[27] Such structural models are often used to organize experimental data in the literature regarding cognitive hierarchy and level-k models.[28] I use such a model here because the predicted behavior for different types matches that of the behavioral types I posit, allowing for noise. This method also allows us to include a fourth *Random* type for robustness.

I hypothesize the existence of four types of players. Players of the lowest level are *random*: they either accept or reject with equal probability. Players of the next level correspond to the *naïve* type. They give a "logit response" (as opposed to best response) to the belief that player 1 is randomizing. Players of the next level correspond to the *sophisticated* type. They give a "logit response" to the belief that player 1 is naïve. Finally, there is an intermediate level player that plays as a naive type in *Sim* and as sophisticated in *Seq*. This is the *GM* type.

For each type the model gives us a distribution of choices conditional on the format and Q. These distributions further depend on a common nuisance parameter λ, which enters the logit response function. If λ goes to infinity, each type's behavior (except for *Random*) converges to that of the prototypical type. If λ goes to zero, all types randomize uniformly. I then use an expectation-maximisation type algorithm to find the value of λ that maximizes the likelihood that the data is explained by our model. At the same time the algorithm classifies subjects into different types.[29] A formal specification of the model is given in the appendix.

I should emphasize that the model here is used purely for the purpose of classification and I do not make any attempt to interpret the model as support for a specific behavioral mechanism that explains the observed behavior. I address this question in the next subsection by leveraging the novel feature of my experiment which is the *PtO* format and comparing behavior there across types. I should note that the general findings are robust to different classification methods as well.[30]

Table 6 summarizes the results of the classification excersise. The vast majority of

[26]These observations are not included in the scatter plot.

[27]See Stahl and Wilson (1994) Stahl and Wilson [1994].

[28]See Crawford et al (2010) Crawford et al. [2010] and references therein.

[29]See Bhatt et al. 2010 Bhatt et al. [2010] for a similar application.

[30]For instance, in previous versions of this manuscript classification followed the use of the k-means clustering algorithm. The downisde of that approach is that it did not allow for the inclusion of a random type.

Naïve	Sophisticated	GM	Random	λ	(st. error)
30	32	30	5	6.676	(0.666)

Table 6: Classification results. The numbers in the first four columns indicate the number of subjects classified as a given type. The last two columns show the estimated value and standard error for the noise parameter μ.

subjects are classified as belonging to one of the three benchmark types in roughly equal proportions. Only 5 subjects are classified as random types. The scatter plot in figure 1 uses different colors to indicate the type assigned to each of the subjects depicted.[31]

4.3 Comparison of types' behavior in PtO

We now turn our attention to subjects' behavior in *PtO*. Recall, that in this format subjects were offered the option to pay 10 points to observe the computer's choice. In other words, subjects could turn the game from simultaneous to sequential by paying this small amount. This choice was made before observing their private signal.

	Naive				Sophisticated				GM			
paid	*No*		*Yes*		*No*		*Yes*		*No*		*Yes*	
	65.7%		**34.3%**		**67.2%**		**32.8%**		**52%**		**48%**	
signal	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*
frequency of reject %	**99.4**	**17.5**	**96.7**	**10**	**100**	**68.2**	**100**	**50**	**100**	**25.6**	**95.2**	**100**
# of obs.	48	149	17	10	58	157	17	4	39	117	26	21

Table 7: Behavior by type in the PtO game. The first row indicates how often each type chooses to pay to observe player 1's choice. The next row shows the frequency of rejection of the barrel, in the *PtO* format.

Recall from table 3 that subjects pay for this option on average 37.7% of the time. Table 7 breaks down this average for each type.[32] While naïve types pay to observe

[31] There are cases of subjects that have very similar, or even the exact same, frequencies of rejection in the two formats but are classified as different types. This happens because the classification algorithm takes into account more information than the one contained in the graph. In particular it takes in to account the value of Q associated with each decision made by a subject.

[32] I do not include the 5 subjects classified as random types in the analysis of this subsection.

slightly more often than sophistcated types, the average for both is much lower than the one for GM types. Note that choices in *PtO* were not use when clasifying individuals into types. Therefore, this difference in behavior for GM types in itself indicates that our classification scheme captures some intrinsic difference between types and is not just an artifact of the assumption that such an additional type exists. This impression is further strengthened by the frequency of rejection for all types in the subsequent decision, to accept or reject the barrel. Again, while this part of the data is not used in order to classify subjects in the previous section, behavior by type conforms to the respective behavioral benchmark.

Hypothesis 3 stated that only *GM* types would pay to observe the computer's choice and the other two types would not. The intuition behind this is that *GM* types, being the only ones behaving differently when the game is played simultaneously vs. sequentially, will be the ones having an incentive to pay to change the format. If that were the case, then, at least for these types it could be argued that the FCT is independent from a failure to perceive the correlation between the computer's choice and the information it received. The data does not support this hypothesis in its strong form. Still, some of the qualitative features of the hypothesis seem to hold. Namely, *Naive* and *Sophisticated* types have similar low rates of paying to observe, and *GM* types do so at a higher rate.

Result 3 *Hypothesis 3* is only qualitatively supported by the data. Individuals of all types pay to observe the computer's choice in a non-negligible number of instances. *GM* types do so about half of the times on average. *Sophisticated* and *Naive* types do so at a lower rate of roughly a third of the times.

An alternative possibility is that the above differences are spurious. This would be consistent with the notion that the FCT for both *Naive* and *GM* types could be explained by a failure to take into account the correlation between the computer's information and choices. I test for no-differences between types below. Before doing so, and to better understand what is happening, it is worth taking a closer look at individuals PtO behavior by type.

Figure 2 presents the average frequency of paying to observe as a function of Q and separated by type. It is clear that the higher average frequency for GM types in the aggregate is driven by a very high frequency for medium values of Q. For higher or lower values, GM types pay to observe as often as other types, and even less for the extreme values of $Q = 1$ or $Q = 9$. No such inverted-U relation appears to exist for the other two types. In fact, the average frequency does not move away much from the average for either of the other two types.

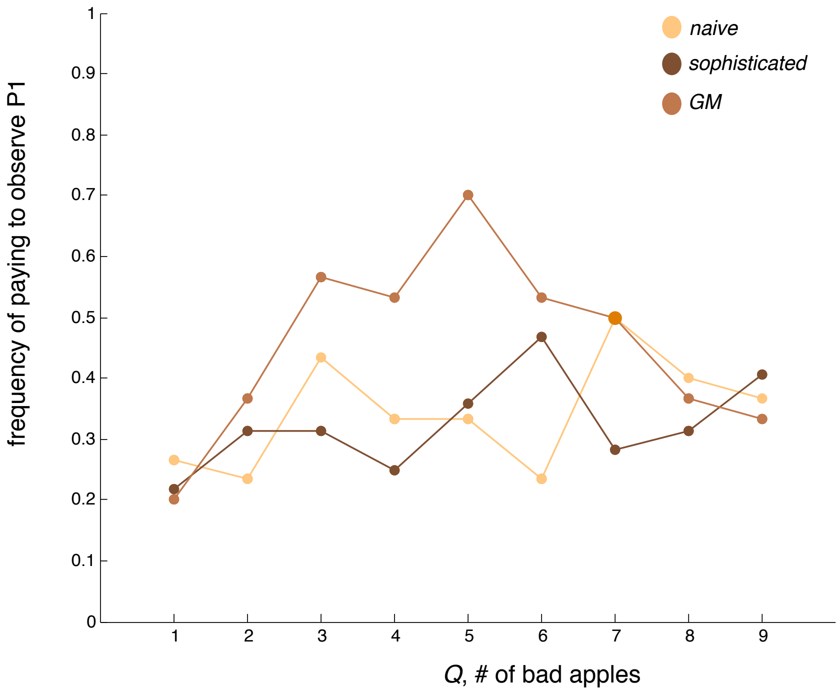

Figure 2: Paying to observe player 1's choice, by type and as a function of Q.

I run separate probit regressions for each type with the decision to pay to observe the computer's choice as a dependent variable. The estimated coefficients are reported in table 8. These results confirm what can be seen in figure 2: the choice to pay to observe does not depend on Q for naive and sophisticated types. The opposite holds true for GM types. The coefficients for both Q and Q^2 are significant and their signs conform to the inverted-U shaped relationship with a maximum at $Q = 5$.

I did not anticipate this pattern of behavior in the *Pto* format when designing the experiment. Nevertheless, it is consistent with the idea that GM types understand that observing the computer's choice reveals its private information. They fail to do the appropriate contingent thinking that would bring the understanding that such information is irrelevant, as it only matters when the computer rejects. This insight is captured by the NCT model introduced in section 2.3, where the decision maker suffers from a pure FCT. As I show in equation 9, for such an individual the value of paying to observe depends non-monotonically on Q and is maximized for $Q = 5$. The VC model, while being able to capture the change in behavior of *GM* types when going from the

Probit regressions
Dependent variable: *paid to observe ('paid' = 1, 'otherwise' = 0)*

	Naive	Sophistcated	GM
Q	0.121	0.130	0.692**
	(0.112)	(0.108)	(0.124)
Q^2	-0.008	-0.009	-0.066**
	(0.112)	(0.011)	(0.012)
round	-0.215**	-0.004	-0.019**
	(0.008)	(0.005)	(0.007)
bad color	0.006	-0.040	0.284
	(0.137)	(0.139)	(0.160)
Intercept	-0.335	-0.727**	-1.223**
	(0.283)	(0.290)	(0.341)
# of obs.	300	320	300

* p < .05 ** p < .01

Note: *Standard errors are clustered at the subject level*

Table 8: Subjects' decision to pay to observe player 1's choice, by type. Explanatory variables are: Q (the number of bad apples in the bad barrel, takes values from 1 to 9); *round* (round in which a given decision is made, takes values from 1 to 40); *bad color* (value 1 if the bad apples/balls were represented by red, 0 if blue).

simultaneous to the sequential format, does not capture the relationship between Q and the choice of paying to observe for these types.

The behavior of *Sophisticated* and *Naive* types is captured nicely by the R and C models respectively (see section 2.3), provided we allow for the possibility of noise in behavior. In what follows I use these models to build and estimate structural models of behavior that can organize our entire dataset. As we shall see, this exercise will support my conjecture about the drivers of the FCT for different behavioral types.

A final note on the results reported in Table 8 is that *round* is also significant and has a negative sign for Naive and GM types, although the magnitude is very small for the latter. This means that the behavior of these types with respect to the decision to pay to observe changes during the experiment in the direction of optimality. It is possible that this is driven by subjects better understanding the situation. Still, such an explanation is hard to reconcile with the fact that they still fail to correct their behavior to the same extend during the experiment when facing the choice of accepting or rejecting the barrel. There is another possibility to consider. In about three out of four times, the computer receives a good signal and accepts the barrel. This means that if the subject pays to observe the computer's choice they will most likely not be given the chance to accept or reject the barrel, which can be frustrating. It is possible that these subjects choose to pay to observe less frequently in order to avoid this frustration. I revisit the issue of learning during the experiment in the analysis of structural models of behavior in the next section.

4.4 Structural analysis

The analysis above led to some interesting findings. First I find that individuals have trouble choosing optimally in the BoA game. Behavior is closer to optimal when the game is played sequentially. Next, I found that there is substantial heterogeneity in behavior in the Accept/Reject choice stage of the game and how it depends on the format. This heterogeneity is largely captured by three behavioral types. Finally, I show that these types could be associated to some underlying differences in cognitive abilities. Naive types do not appear to be able to grasp the correlation between player 1's information and her choice. On the other hand, GM types seem to recognize the existence of such correlation, but fail to correctly compute the optimal choice when they need to reason about it contingently.

To assess the robustness of our previous conclusions I perform here additional analysis that does not rely on the classification of individual subjects. This allows me to include all individuals, even if it was not possible to classify their behavior into a specific behavioral type. I estimate structural models of behavior that give predictions for the entire set of choices that individuals face in the experiment: both Accept/Reject and pay to observe choices.

Mixture models

Models included	Model weights				Nuisance	log Likelihood
	$\hat{\gamma}_R$	$\hat{\gamma}_C$	$\hat{\gamma}_{VC}$	$\hat{\gamma}_{NTC}$	$\hat{\lambda}$	
R	1	-	-	-	1.93 (.101)	- 1910.2
C	-	1	-	-	3.56 (.175)	- 1735.3
VC	-	-	1	-	2.56 (.123)	- 1794.6
NCT	-	-	-	1	2.09 (.088)	- 1743.2
R, C	0.33 (.052)	0.67	-	-	4.96 (0.249)	- 1612.1
C, VC	-	0.54 (.059)	0.46	-	4.87 (0.247)	- 1608.0
C, NCT	-	0.63 (.057)	-	0.37	4.16 (0.213)	- 1595.8
R, C, VC	0.16 (.042)	0.52 (.056)	0.32	-	6.00 (0.307)	- 1550.2
R, C, NCT	0.21 (.045)	0.51 (.055)	-	0.28	5.85 (0.309)	- 1500.9
R, C, VC, NCT	0.19 (.045)	0.50 (.055)	0.05 (.045)	0.26	6.04 (0.341)	- 1499.9

R = *Rational expectations*, C = *Cursed*, VC = *Varying Cursedness*, NCT = *Non Contingent Thinker*

Note: *Standard errors in parentheses. Calculated using the inverse of the Hessian matrix.*

Table 9: Structural estimation of a selection of mixture models. Each model is fitted on the complete data set including Accept/Reject choices under all three formats, as well as 'pay to observe' choices for the *PtO* format. For type weights I estimate the weight for $k-1$ types and assign the remaining weight to the last type. Therefore, each model with k types has k free parameters: the $k-1$ weights plus the nuissance parameter λ.

The models of cognitive reasoning I use in this analysis are based on the ones introduced in Section 2.3: Rational expectations, Cursed, Varying Cursedness, Non Contingent Thinker. In all cases I allow for "mistakes" by having individuals making logit-choices with a nuissance parameter λ. While I do fit each individual model separately to the data, to address the observed heterogeneity I also fit mixture models that allow different weighted combinations of the individual models. That is, I estimate the weight γ_i for each model i in a subset of the model space and restrict the weight of the remaining models to be zero. Table 9 shows the results of this exercise for a selection of the estimated models.[33]

I compare the fit of the different non-nested models to the data using the Vuong closeness test ([Vuong, 1989]). The mixture of R, C and NCT (2nd to last row of Table 9) fits the data better than all mixtures of three or fewer models.[34] Compared to that, the improvement in fit from the mixture of all four models (last row) is not statistically significant.[35] Overall, the results support findings from our previous analysis. Both the C and NCT models add substantial explanatory power to the mixture models that incude them. Recall that in both those models decison makers do not behave optimally when the game is played simultaneously. In the NCT model the decision maker does play optimally in the sequential format. While the VC model also has this feature, it does not predict any difference in behavior in the first stage of the PtO format compared to the R or C models. The explanatory power of the NCT model indicates that there is significant correlation between how behavior depends on Q in the first stage of the PtO format and how subjects' behavior in the Acept/reject stage depends on whether the game is played simultaneously or sequentially.

As the previous analysis indicated that some learning might be taking place, I also consider an extendion of these mixture models where I introduce a linear time trend.[36] A positive value of the time trend parameter δ would indicate that the optimal choices (i.e. not paying to observe and rejecting the barrel) become more attractive as rounds go by. Given the estimation results, this does not seem to be the case. If anything, δ is estimated to be negative for most mixture models. Alternative ways of specifying the time trend give very similar results.[37] Given the results presented in Tables 4 and 8, my

[33] Here I only present the best fitting mixtures from every two- and three-model mixtures. The complete results are available upon request.

[34] In all comparisons we can reject the null of the models being equally close to the true data-generating-process based on a Vuong closeness test ($p - val < .01$ for all comparisons).

[35] We cannot reject the null of the models being equally close to the true data-generating-process when using the Vuong closeness test ($p - val = .77$).

[36] See Appendix B.

[37] These are available upon request.

interpretations of these findings is that while learning effects cannot be ruled out, these are unlikely to be driving any of the main results reported.

5 Conclusions

This paper adds to the effort to understand individuals' susceptibility to the winner's curse, and how this is connected to failures of contingent thinking. To this end I design and run a lab experiment using a simple game that incorporates a winner's curse environment. I find substantial heterogeneity in how subjects in our experiment react to differences in the game's format. One subset is reasonably good in avoiding the curse independently of the games. These sophisticated subjects seem to perform both computations well. A second subset of similar size suffers from the curse in all formats. This naïve behavior is mainly the result of an inability to account for the correlation between states of the world and their opponents actions. Even when they observe the computer rejecting the barrel, they are unable to infer that the state must be bad. The difference in aggregate behavior across game formats is almost entirely attributed to a third subset of subjects. Our design leads us to conclude that these individuals have no problem in understanding how their opponents actions reveal their private information, neither in the simultaneous or the sequential game formats. They are not "cursed" (nor suffer from correlation neglect). Their behavior is purely the result of a failure in performing the necessary contingent thinking.

Heterogeneity in individuals' behavior is widely observed in economic experiments, which perhaps explains the popularity of models of strategic thinking that incorporate such heterogeneity (e.g. Level-k). Our experiment uncovers heterogeneity of a higher order: individuals not only differ in their behavior, but also in how this behavior is affected by the strategic environment. This poses new challenges for decision and game theorists that want to develop parsimonious models to capture behavior in different settings. Such models are necessary for economists studying the design of markets and other institutions, where issues such as who can observe whom and when are critical.

On the bright side, I present evidence that the underlying differences in cognitive abilities and modes of reasoning that drive this heterogeneity are not affected by the environment in which decisions take place. Decision neuroscience has recently made rapid developments and could perhaps allow us to better understand the nature of the cognitive calculations done during decision making as well as the potentialy different abilities of individuals at that level. Anyone with a desire to accurately model strategic behavior should keep a vigilant eye on that research and be ready to incorporate such findings into formal models.

References

Marina Agranov, Elizabeth Potamites, Andrew Schotter, and Chloe Tergiman. Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 2012.

Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *The Review of Economic Studies*, 83(4):1297–1333, 2016.

Larbi Alaoui and Antonio Penta. Cost-benefit analysis in reasoning. *Journal of Political Economy*, 130(4):881–925, 2022.

S. Nageeb Ali, Jacob K. Goeree, Navin Kartik, and Thomas R. Palfrey. Information aggregation in standing and ad hoc committees. *The American Economic Review*, 2008.

Felipe A Araujo, Stephanie W Wang, and Alistair J Wilson. The times they are a-changing: dynamic adverse selection in the laboratory. *American Economic Journal: Microeconomics*, 13(4):1–22, 2021.

Max H. Bazerman and William F. Samuelson. I Won the Auction But Don't Want the Prize. *Journal of Conflict Resolution*, 1983.

Meghana A. Bhatt, Terry Lohrenz, Colin F. Camerer, and P. Read Montague. Neural signatures of strategic types in a two-person bargaining game. *Proceedings of the National Academy of Sciences*, 2010.

Barry Blecherman and Colin F Camerer. Is there a winner's curse in the market for baseball players? evidence from the field. *Social Science Working Paper - Caltech*, 1996.

Colin F. Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 2004.

Edward C. Capen, Robert V. Clapp, and William M. Campbell. Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 1971.

Timothy N Cason and Charles R Plott. Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy*, 122(6):1235–1270, 2014.

James Cassing and Richard W Douglas. Implications of the auction mechanism in baseball's free agent draft. *Southern Economic Journal*, pages 110–121, 1980.

Gary Charness and Dan Levin. The Origin of the Winner's Curse: A Laboratory Study. *American Economic Journal: Microeconomics*, 2009.

Shani Cohen and Shengwu Li. Sequential cursed equilibrium. *arXiv preprint arXiv:2212.06025*, 2022.

Vincent P. Crawford and Nagore Iriberri. Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions? *Econometrica*, 2007.

Vincent P. Crawford, Miguel A. Costa-Gomes, and Nagore Iriberri. Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications. *Journal of Economic Literature*, 51(1):5–62, 2010.

Benjamin Enke and Florian Zimmermann. Correlation neglect in belief formation. *The review of economic studies*, 86(1):313–332, 2019.

Ignacio Esponda and Emanuel Vespa. Hypothetical Thinking and Information Extraction: Strategic Voting in the Laboratory. *American Economic Journal: Microeconomics*, 2014.

Ignacio Esponda and Emanuel Vespa. Endogenous sample selection: A laboratory study. *Quantitative Economics*, 9(1):183–216, 2018.

Erik Eyster and Matthew Rabin. Cursed Equilibrium. *Econometrica*, 2005.

Erik Eyster and Georg Weizsacker. Correlation neglect in financial decision-making. 2010.

Urs Fischbacher. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 2007.

Meng-Jhang Fong, Po-Hsuan Lin, and Thomas R Palfrey. Cursed sequential equilibrium. *arXiv preprint arXiv:2301.11971*, 2023.

Daniel Friedman. Monty hall's three doors: Construction and deconstruction of a choice anomaly. *The American Economic Review*, 88(4):933–946, 1998.

Uri Gneezy. Step-level reasoning and bidding in auctions. *Management Science*, 2005.

Ben Greiner. An online recruitment system for economic experiments. *Working Paper Series in Economics 10*, 2004.

Kenneth Hendricks, Robert H Porter, and Bryan Boudreau. Information, returns, and bidding behavior in ocs auctions: 1954-1969. *The Journal of Industrial Economics*, pages 517–542, 1987.

Asen Ivanov, Dan Levin, and Muriel Niederle. Can Relaxation Of Beliefs Rationalize The Winner's Curse?: An Experimental Study. *Econometrica*, 2010.

Philippe Jehiel. Analogy-based expectation equilibrium. *Journal of Economic Theory*, 2005.

Philippe Jehiel and Frédéric Koessler. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 2008.

John H Kagel and Dan Levin. The winner's curse and public information in common value auctions. *American economic review*, 76(5):894–920, 1986.

John H. Kagel and Dan Levin. *Common value auctions and the winner's curse*. Princeton Univeristy Press, 2002.

John H Kagel, Ronald M Harstad, Dan Levin, et al. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55 (6):1275–1304, 1987.

John H Kagel, Dan Levin, and Ronald M Harstad. Comparative static effects of number of bidders and public information on behavior in second-price common value auctions. *International Journal of Game Theory*, 24(3):293–319, 1995.

Mario Levis. The winner's curse problem, interest costs and the underpricing of initial public offerings. *The Economic Journal*, 100(399):76–89, 1990.

Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11): 3257–3287, 2017.

Philippos Louis. *Information, behavior, and the design of institutions*. PhD thesis, Universitat Autònoma de Barcelona, Departament d'Economia i d'Història Economica, 2012.

Scott E Maxwell and Harold D Delaney. Designing experiments and analyzing data: A model comparison perspective. 2003.

Walter J Mead, Asbjorn Moseidjord, and Philip E Sorensen. The rate of return earned by lessees under cash bonus bidding for ocs oil and gas leases. *The Energy Journal*, 4 (4):37–52, 1983.

Johannes Moser. Hypothetical thinking and the winner's curse: An experimental investigation. *Theory and Decision*, 87(1):17–56, 2019.

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, 1995.

M Kathleen Ngangoué and Georg Weizsäcker. Learning from unrealized versus realized prices. *American Economic Journal: Microeconomics*, 13(2):174–201, 2021.

Muriel Niederle and Emanuel Vespa. Cognitive limitations: Failures of contingent thinking. *Annual Review of Economics*, 15(1):307–328, 2023.

Hyoeun Park. Complexity and contingent reasoning. *Experimental Economics*, 28(1): 200–215, 2025.

Evan Piermont and Peio Zuazo-Garin. Failures of contingent thinking. *arXiv preprint arXiv:2007.07703*, 2020.

Matthew Rabin and Georg Weizsäcker. Narrow bracketing and dominated choices. *American Economic Review*, 99(4):1508–1543, 2009.

Kevin Rock. Why new issues are underpriced. *Journal of financial economics*, 15(1-2): 187–212, 1986.

Richard Roll. The hubris hypothesis of corporate takeovers. *Journal of business*, pages

197–216, 1986.
William Samuelson and Max H Bazerman. The winner's curse in bilateral negotiations. 1984.
Eldar Shafir and Amos Tversky. Thinking through uncertainty: Nonconsequential reasoning and choice. *Cognitive psychology*, 24(4):449–474, 1992.
Dale O. Stahl and Paul W. Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 1994.
Dale O. Stahl and Paul W. Wilson. On Players Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 1995.
Richard Thaler. Anomalies: The winner's curse. *The Journal of Economic Perspectives*, 1988.
Edward F Vonesh. Efficiency of repeated measures designs yersus completely randomized designs based on multiple comparisons. *Communications in Statistics-Theory and Methods*, 12(3):289–301, 1983.
Quang H Vuong. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: journal of the Econometric Society*, pages 307–333, 1989.
Roberto A. Weber. 'Learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 2003.

A Classification

A.1 A hierarchy of logit choice models

First, let $f_i : X^2 \to X$ be the assignment function for player i. In particular,

$$f_1(x_1, x_2) = x_1$$

and

$$f_2(x_2, x_1) = \begin{cases} x_2, & for\ x_1 = R \\ R, & for\ x_1 = A \end{cases}$$

For notational economy I use the set of choices to also denote the set of outcomes. The outcome A means the player keeps the object, while the outcome R means he does not keep it. Payoffs are given by

$$u_i\left(f_i(x_i, x_{-i}), \theta\right) = \begin{cases} 1, & f_i(x_i, x_{-i}) = A \ \& \ \theta = G \\ 1, & f_i(x_i, x_{-i}) = R \ \& \ \theta = B \\ 0, & otherwise \end{cases}$$

Let $\rho^k_{r,s}$ be the probability of rejecting for a subject of type $k \in \{naive\,,\ soph\,,\ GM\,,\ random\}$ in format $r \in \{sim, seq\}$ after observing signal $s \in \{g, b\}$.

A *naïve* type, incorrectly believes that the computer makes a uniformly random choice between accepting and rejecting. Player 2 then makes a logit response with a noise parameter λ. The probability of rejecting for this type is:

$$\begin{aligned} \rho^{naive}_{sim,s}(q,\lambda) = & \left(1 - \frac{1}{exp\{-\lambda\left[u^e_2\left(f_2(R,x_1),\theta|s\right) - u^e_2\left(f_2(A,x_1),\theta|s\right)\right]\}}\right)^{-1} & ,\forall s \\ \rho^{naive}_{seq,s}(q,\lambda) = & \left(1 - \frac{1}{exp\{-\lambda\left[u^e_2\left(f_2(R,R),\theta|s\right) - u^e_2\left(f_2(A,R),\theta|s\right)\right]\}}\right)^{-1} & ,\forall s \end{aligned}$$

The *sophisticated* type has correct but noisy beliefs about the computer's choice, and makes a noisy response to such beliefs. The probability of rejecting for the *sophisticated* type is:

$$\begin{aligned} \rho^{soph}_{sim,s}(q,\lambda) = & \left(1 - \frac{1}{exp\{-\lambda\left[u^e_2\left(f_2(R,x_1),\theta,\lambda|s\right) - u^e_2\left(f_2(A,x_1),\theta,\lambda|s\right)\right]\}}\right)^{-1} & ,\forall s \\ \rho^{soph}_{seq,s}(q,\lambda) = & \left(1 - \frac{1}{exp\{-\lambda\left[u^e_2\left(f_2(R,R),\theta,\lambda|s\right) - u^e_2\left(f_2(A,R),\theta,\lambda|s\right)\right]\}}\right)^{-1} & ,\forall s \end{aligned}$$

The *GM* type behaves like the *naïve* type in the case of simultaneous decisions and as the *sophisticated* type in the case of sequential decisions:

$$\begin{aligned} \rho^{GM}_{sim,s}(q,\lambda) = & \ \rho^{naive}_{sim,s}(q,\lambda) & ,\forall s \\ \rho^{GM}_{seq,s}(q,\lambda) = & \ \rho^{soph}_{seq,s}(q,\lambda) & ,\forall s \end{aligned}$$

Finally, I allow for a *random* type, whose decision is a uniformly random draw between 'accept' and 'reject':

$$\rho_{r,s}^{rand} = \frac{1}{2} \quad , \forall r\,,\ s$$

A.2 Maximum likelihood estimation and classification.

The noise parameter $\lambda > 0$ is estimated using maximum likelihood estimation. Classification of subjects to different levels is done within the MLE routine. Let $C_i = \{(x_i^t, r_i^t, s_i^t) : x_i^t \in X\,,\ r_i^t \in \{sim, seq\}\,,\ s_i^t \in \{g, b\}\,,\ t \in T_i\}$ represent the choices made by subject $i \in N = \{1, \dots, 97\}$ in round t where the format for the subject was r_i^t and he observed a signal s_i^t. T_i is the set of rounds in which subject i had to make decisions. Recall that in format *Seq*, subjects take no decision if the computer accepts. Let $L^{k_i}(C_i, \lambda)$ be the interim likelihood function for the choices of subject i classified as level k_i, that made choices C_i. In particular it is:

$$L^{k_i}(C_i, \lambda) = \Pi_{t \in T_i} \left((1 - \rho^{k_i}_{r_i^t, s_i^t}) x_i^t + \rho^{k_i}_{r_i^t, s_i^t} (1 - x_i^t) \right)$$

The likelihood function maximized in MLE is

$$L(x, \lambda) = \Pi_{i \in N} \max_{k_i \in \{0,1,2\}} \{L^{k_i}(x_i, \lambda)\}$$

For computational reasons, as is standard, the logarithm of the likelihood function was used.

To calculate the CI's for the frequencies in each case (format-signal combination) I use Monte Carlo simulations. I simulate the model 2000 times using the estimated value for λ, the classification obtained and the the actual set of choices made by each subject. This gives a distribution of aggregate rejection frequencies for each of the four cases (format & private signal): simultaneous - bad signal, simultaneous good signal, sequential - bad signal, sequential - good signal. The lower and upper bounds of the 95% CI for each case are the 2.5 and 97.5 percentiles of the corresponding distribution.

B Structural analysis details

Let $\rho_{r,s}^m$ be the probability of rejecting for a subject following the model of cognitive reasoning $m \in \{R\,,\ C\,,\ VC\,,\ NCT\}$ in format $r \in \{sim, seq\}$ after observing signal $s \in \{g, b\}$.

Assuming logit responses, it is given by:

$$\rho^{m}_{sim,s}(q,\lambda) = \quad \left(1 - \frac{1}{exp\{-\lambda\left[U^r_m(R|\cdot) - U^r_m(A|\cdot)\right]\}}\right)^{-1} \quad , \forall s, m, \tag{10}$$

where U^r_m are given by expressions (1)-(8) in Section 2 and λ is a nuisance parameter.

Also, let π^m be the probability of *paying to observe* for a subject following the model of cognitive reasoning $m \in \{R,\ C,\ VC,\ NCT\}$ in format $r \in \{sim, seq\}$ after observing signal $s \in \{g, b\}$. Following the analysis in Section 2, and assuming logit responses with the same nuisance parameter λ, this is given by:

$$\pi^m(q,\lambda) = \quad \left(1 - \frac{1}{exp\{\lambda \cdot c\}}\right)^{-1} \quad , \forall m \in \{R, C, VC\} \quad , \tag{11}$$

and

$$\pi^{NCT}(q,\lambda) = \quad \left(1 - \frac{1}{exp\{-\lambda\left[\frac{V}{2}q(1-q)\right]\}}\right)^{-1} \quad , \tag{12}$$

where the last expression follows from equation (9) and c represents the cost a subject needs to pay to observe player 1's choice.

We represent the choice whether to pay to observe in the *PtO* format as $y \in Y = \{pay, no\}$. Let $C_{i,t} = \{(x, y, r, s, Q) : x \in X \cup \{\varnothing\},\ \in Y \cup \{\varnothing\},\ r \in \{sim, seq\},\ s \in \{g, b\},\ q \in \{.1, ..., .9\}\}$ represent the choices x and y made by subject $i \in N = \{1, \ldots, 97\}$ in round $t \in T = \{1, ...40\}$ where the format for the subject was r, they observed a signal s, and there were $q \cdot 10$ bad apples in the bad barrel. Then $C = \{\mathrm{C}_{i,t}\}_{\forall i \in N, t \in T}$ is the collection of all Let γ_m represent the fraction of the population that follows the cognitive reasoning model $m \in M = \{R, C, VC, NCT\}$. Let $I_(\cdot)$ represent the indicator function, taking values in $\{0, 1\}$. Then the likelihood function is given by:

$$\begin{aligned} L(C,\lambda) = \quad & \Pi_{i\in N}\Pi_{t\in T_i} \cdot \textstyle\sum_{m\in M} \gamma_m \quad \Bigg[\left(\rho^m_{r^t_i, s^t_i}(q^t_i,\lambda) \cdot I(x^t_i = A) + (1 - \rho^m_{r^t_i, s^t_i}(q^t_i,\lambda)) \cdot I(x^t_i = R)\right)^{I(x^t_i \neq \varnothing)} \\ & + \left(\pi^m(q^t_i,\lambda) \cdot I(y^t_i = pay) + (1 - \pi^m(q^t_i,\lambda) \cdot I(y^t_i = no)\right)^{I(y^t_i \neq \varnothing)}\Bigg] \end{aligned} \tag{14}$$

The parameter values that maximize the logarithm of this function are reported in Table 9.

We repeat the same exercise introducing a "time trend" as follows. In expressions (10)-(12) we introduce a a linear component $\delta \cdot \frac{t}{40}$. For $\delta > 0$ this represents the otpimal choice becoming more attractive in later rounds. In particular we have:

$$\rho^{m}_{sim,s}(q,\lambda) = \left(1 - \frac{1}{exp\left\{-\lambda\left[U^{r}_{m}(R|\cdot) - U^{r}_{m}(A|\cdot) + \delta\cdot\frac{t}{40}\right]\right\}}\right)^{-1} \quad, \forall s, m,$$

$$\pi^{m}(q,\lambda) = \left(1 - \frac{1}{exp\left\{-\lambda\left[-\right]c + \delta\cdot\frac{t}{40}\right]\right\}}\right)^{-1} \quad, \forall m \in \{R, C, VC\} \quad ,$$

and

$$\pi^{NCT}(q,\lambda) = \left(1 - \frac{1}{exp\left\{-\lambda\left[\frac{V}{2}q(1-q) + \delta\cdot\frac{t}{40}\right]\right\}}\right)^{-1} \quad ,$$

The estimated parameter values for these models are shown in Table 10.

Mixture models with Time trend

Models included	Model weights				Nuisance	Time trend	log Likelihood
	$\hat{\gamma}_R$	$\hat{\gamma}_C$	$\hat{\gamma}_{VC}$	$\hat{\gamma}_{NTC}$	$\hat{\lambda}$	$\hat{\delta}$	
R	1	-	-	-	2.38	-0.30	- 1861.1
					(.114)	(.028)	
C	-	1	-	-	3.57	.004	- 1735.2
					(.180)	(.019)	
VC	-	-	1	-	2.61	-0.014	- 1779.5
					(.119)	(.026)	
NCT	-	-	-	1	2.11	-0.013	- 1735.2
					(.087)	(.033)	
R, C	0.37	0.63	-	-	4.80	-0.075	- 1601.3
	(.054)				(0.234)	(0.017)	
C, VC	-	0.53	0.47	-	4.76	-0.021	- 1607.2
		(.059)			(0.252)	(0.016)	
C, NCT	-	0.63	-	0.37	4.10	-0.018	- 1595.3
		(.057)			(0.217)	(0.019)	
R, C, VC	0.17	0.51	0.32	-	5.77	-0.033	- 1547.5
	(.044)	(.056)			(0.306)	(0.015)	
R, C, NCT	0.22	0.50	-	0.28	5.63	-0.042	- 1496.7
	(.046)	(.055)			(0.303)	(0.015)	
R, C, VC, NCT	0.21	0.50	0.02	0.27	5.72	-0.039	- 1496.6
	(.048)	(.055)	(.046)		(0.346)	(0.016)	

R = *Rational expectations*, C = *Cursed*, VC = *Varying Cursedness*, NCT = *Non Contingent Thinker*

Note: *Standard errors in parentheses. Calculated using the inverse of the Hessian matrix.*

Table 10: Structural estimation of a selection of mixture models with a linear time trend. Each model is fitted on the complete data set including Accept/Reject choices under all three formats, as well as 'pay to observe' choices for the *PtO* format. For type weights I estimate the weight for $k-1$ types and assign the remaining weight to the last type. Therefore, each model with k types has k free parameters: the $k-1$ weights plus the nuissance parameter λ.

C Instructions

Thank you for participating in this session. The experiment will involve a series of decisions. Each of you may earn different amounts. The amount you earn depends on your decisions and chance. The exchange rate used in the experiment is 20 points for 1 CHF. You also receive a 10 CHF participation fee. Upon completion of the experiment, you will be paid individually and privately.

Please remain quiet!

You will be using the computer terminal for the entire experiment, and your decisions will be made on your computer terminals. Please DO NOT talk or make any other audible noises during the experiment. If you have any questions, raise your hand and your question will be answered so that everyone can hear.

The experiment consists of a series of offers made to you and another party. The other party will actually be simulated by a computer. Therefore, from now on we refer to it as "the computer". These offers all share a common structure which is the following: You and the computer are offered a single box which you may accept or reject. Since there is only one box that is offered, it is not possible for both you and the computer to get it. In particular, the computer has priority. This means that:

- if you both accept the box, the computer gets it.
- If you both reject it, neither you nor the computer gets it.
- If one accepts and the other rejects, the box is given to the one that accepted.

Payoffs: The box may be GOOD or BAD with equal probability. Your payoffs depend on whether the box is good or bad and whether you get it:

- If you get the box and it is GOOD, you win 100 points.
- If you get the box and it is BAD, you win 0 points.
- If you do not get the box and it is GOOD you win 0 points.
- If you do not get the box and it is BAD you win 100 points.

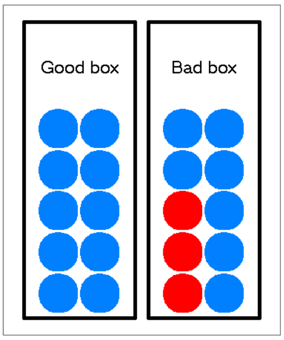

Figure 1

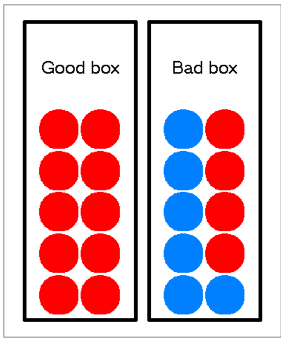

Figure 2

Information: You will not know whether the box is good or bad. You know the following: The box contains 10 balls that are either red or blue. If it is GOOD, all balls have the same color. If it is BAD, it contains both red and blue balls. The following examples show you what information you will have for each offer:

Example 1: In figure 1 you see an example where:

- If the box is good, all balls are blue. (left half)
- If the box is bad, there are 3 red balls and 7 blue balls. (right half)

Example 2: In figure 2 you see an example where:

- If the box is good, all balls are red. (left half)
- If the box is bad, there are 6 blue balls and 4 red balls. (right half)

In each offer you will be shown a similar figure.

Before you and the computer make the decision to accept or reject the box the following happens: the computer draws a ball from the box, looks at its color and returns it to the box. Then, you do the same: you draw a ball, look at its color and return it to the box. These actions will take place virtually and you will see on the screen the color of the ball you drew. You cannot see the color of the ball the computer draws. Neither can the computer see the color of the ball you draw.

The computer's decision: We now explain how the computer makes decisions and then give some examples to make it clear. The computer makes decisions according to the following rule:

If the color of the ball it draws from the box is the same as the color of the balls in the GOOD box, then the computer accepts the box. Otherwise it rejects.

Example 1: If the two possible contents of the box are like the ones depicted in figure 1, then the computer:

- Accepts if he draws a blue ball
- Rejects if he draws a red ball.

Example 2: If the two possible contents of the box are like the ones depicted in figure 2, then the computer:

- Accepts if he draws a red ball
- Rejects if he draws a blue ball.

Situations: There are three different situations in which you will have to make your decisions. You will be put in each of these situations a number of times and in random order:

- **Situation 1.** You and the computer decide simultaneously without observing each other's decisions. (10 times)
- **Situation 2.** The computer decides first and you can observe his decision. If the computer accepts, you cannot do anything. If it rejects, you have to decide yourself. (20 times)
- **Situation 3.** You are given the option to pay 10 points in order to observe the computer?s decision before you decide for yourself. (10 times)

The following screenshots demonstrate what the screen will look like in each offer in these different situations. They all share some common features:

- In the upper left frame you receive a short description of the situation you are in. This corresponds to the three situations described previously.
- On the upper right frame you are shown the two possible contents of the box. If the box is good it will contain the balls shown on the left half of the frame. If the box is bad it will contain the balls shown on the right half. The exact contents will be different in every offer.
- The lower left frame is reserved for information about the computer's decisions. In situation 1 this information is not available, as shown in Screenshot 1. In situation 2 you will be shown the computer's decision (Screenshot 2). In situation 3, this is the place where you make your choice whether or not to buy the option to observe the computer's decision (Screenshot 3).
- In the lower center you can see the ball you have drawn from the box. (Screenshots 1 and 2). In situation 3, this will be shown to you only after you decide whether or not to buy the option to observe the computer's decision (Screenshot 3).

- In the lower right frame you are asked to introduce your decision on whether to accept or reject the box (Screenshot 1).

continues in next page

Screenshot 1: situation 1

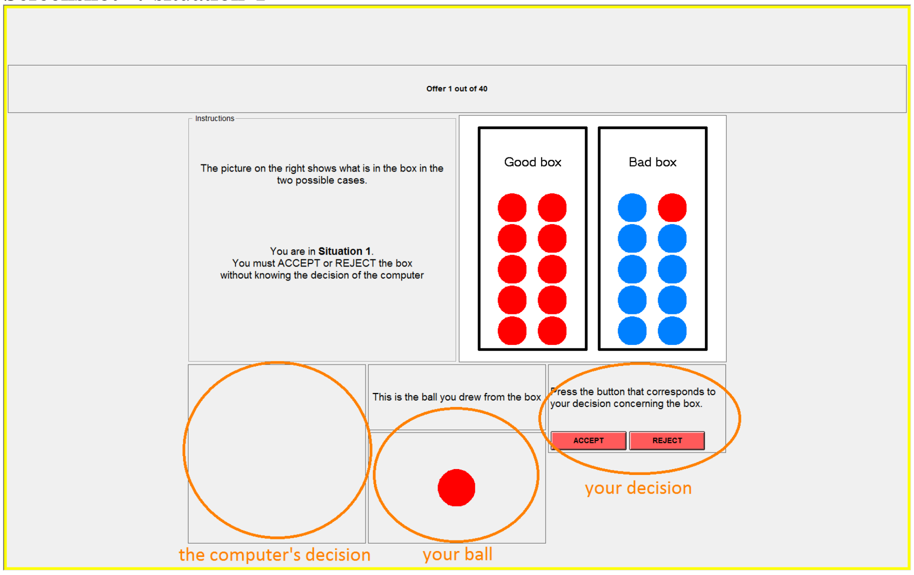

Screenshot 2: situation 2

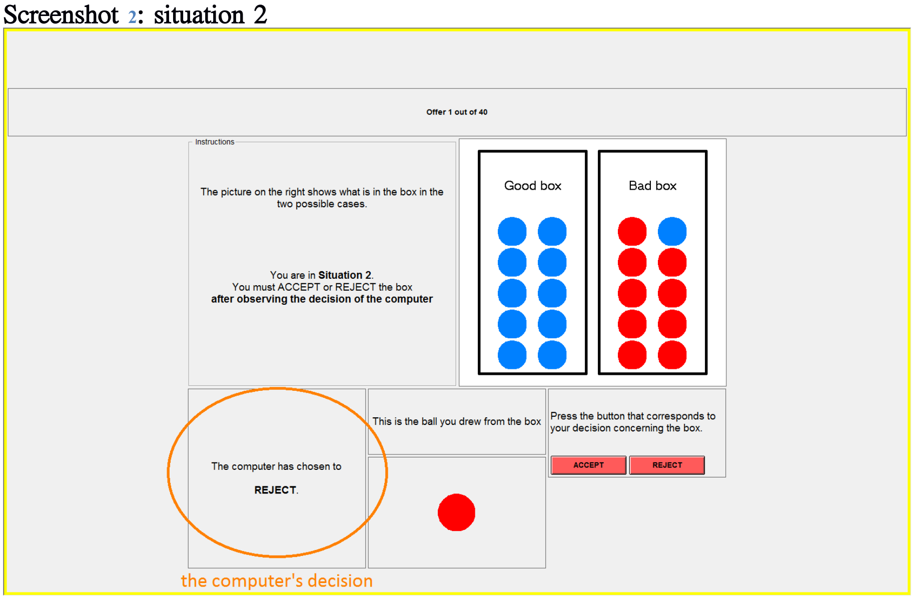

Screenshot 3: situation 3

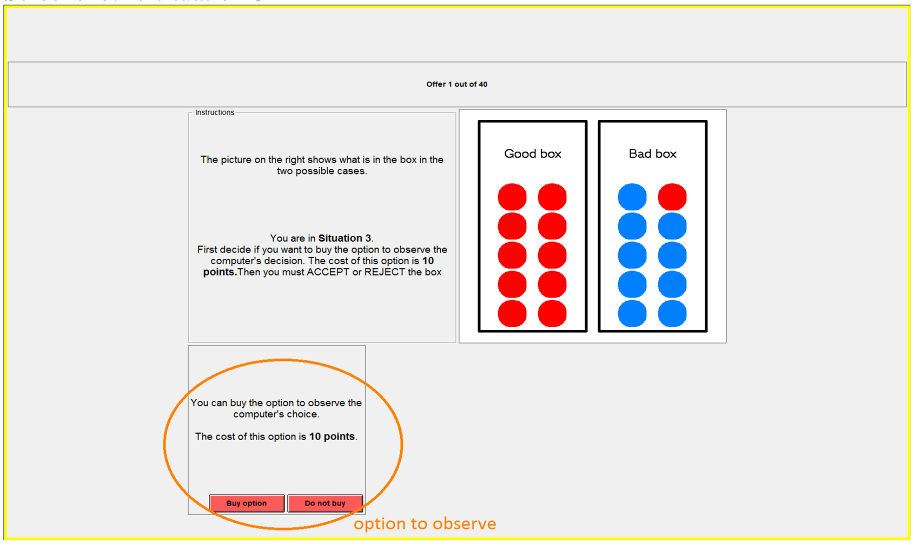

In total you will face 40 different offers. In each of these you get different payoffs depending on your decisions, the computer's decisions, and chance. Here are some examples

Example 1: You are in situation 3. You pay 10 points to observe the computer?s choice. The computer accepts and gets the box. The box turns out to be bad, so you win 100 points because it is bad and you didn't get it. Your payoff in this offer is: -10+100=90 points.

Example 2: You are in situation 2. The computer rejects the box. You accept the box. It turns out to be bad, so you don't win any points because it is bad and you got it. Your payoff in this offer is: 0 points.

Example 3: You are in situation 3. You pay 10 points to observe the computer's choice. The computer rejects the box. You accept the box. It turns out to be bad, so you win 0 points because it is bad and you got it. Your payoff is: -10+0 = -10 points.

To determine your final earnings, 7 of these offers will be chosen randomly and your payoff in these 7 offers will be used to determine your final payoff. At the end you will get an additional 100 points. This is to guarantee that you will have no negative earnings at the end of the experiment.

Quiz: Before starting the sequence of offers you are asked to take a small quiz to make sure you have understood these instructions. You can try each question as many times as necessary until you get it right. If you have any questions during the quiz, raise your hand and one of the assistants will approach you.

Once you finish, wait silently at your seat until we call you to come and receive your payment.

D Comprehension Quiz

After reading the instructions and before starting the experiment, subjects had to answer the following questions.

1. If both you and the computer accept the box:

   a. You keep it and you win 100 points if it is good.
   b. The computer keeps it and you win 100 points if it is good.
   c. You keep it and you win 100 points if it is bad.
   d. The computer keeps it and you win 100 points if it is bad.

2. If you accept the box and the computer rejects:

   a. You keep it and you win 100 points if it is bad.
   b. The computer keeps it and you win 100 points if it is good.
   c. You keep it and you win 100 points if it is good.
   d. The computer keeps it and you win 100 points if it is bad.

3. If you both reject the box:

   a. Nobody keeps it and you win 100 points if it is bad.
   b. You keep it and you win 100 points if it is good.
   c. Nobody keeps it and you win 100 points if it is good.
   d. The computer keeps it and you win 100 points if it is bad.

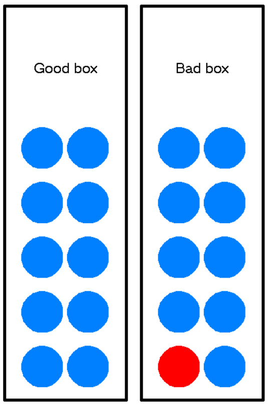

4. Suppose the box is one of the two shown above. Also, suppose that a single ball is drawn and shown to you and the computer (*note: this is not how it is done in the experiment*). If the ball is BLUE, then:

   a. You know that the box is good.
   b. You know that the box is bad.
   c. You know it is more likely for the box to be good.
   d. You know it is more likely for the box to be bad.

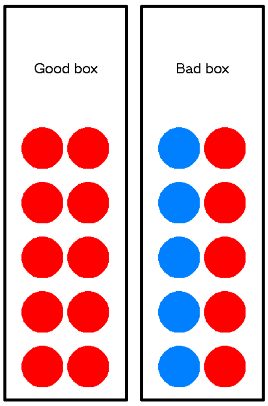

5. Suppose the box is one of the two shown above. Also, suppose that a single ball is drawn and shown to you and the computer (*note: this is not how it is done in the experiment*). If the ball is BLUE, then:

   a. You know that the box is good.
   b. You know that the box is bad.
   c. You do not know for sure whether the box is good or bad.
   d. You know it is more likely for the box to be bad.