The Barrel of Apples Game: contingent thinking, inference from others' actions, and strategic heterogeneity

October 31, 2015

Abstract

In the barrel of apples game players face a winner's curse. It can be played simultaneously or sequentially without affecting behavior, according to common models of strategic thinking. Subjects in an experiment play better in the sequential format. Behavior is heterogeneous within and across formats. We classify subjects in to types based on behavior in the simple game. These types' behavior in an extended version of the game suggests that heterogeneity can be associated to differences in specific cognitive abilities. A failure to infer others' information from their actions explains the winner's curse only for the most naïve subjects. Subjects that behave differently depending on the game's format are able to make such inferences even if they cannot directly observe the other player's action. Still, they fail to perform the necessary contingent thinking to avoid the curse.

Keywords: *contingent thinking, inference from others' actions, cursed equilibrium.*
JEL codes: C72, D82, D83

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

1 Introduction

The preceding quote demonstrates a "winner's curse". Obtaining a membership reveals to Groucho Marx that it actually is undesirable: winning is bad news! The comedian realizes that and resigns. A natural question for a game theorist is: why did he apply in the first place? This paper attempts to provide an answer.

The "winner's curse" is present in a broad range of economic situations, including common-value auctions, financial markets and bargaining. Furthermore, the reasoning necessary to avoid the curse can also be applied in other strategic environments where agents' decisions rely on private information that is relevant to others. Understanding when individuals are able to perform the relevant computations and how this ability is affected by the strategic environment is crucial for the design of markets and other institutions. More generally, it is important if we want to develop better models of strategic thinking that take in to account individuals' differential reasoning abilities in different situations. We use a novel experiment to detect which cognitive computation failures prevent individuals from avoiding the curse in simultaneous and sequential game situations.

There are two cognitive computations that appear to be relevant in"winner's curse" situations and that have received attention in the literature. The first is *Contingent Thinking*, which can be defined as the ability to think as if a state has been realized. The second is *Inference from Others' Actions*, which refers to the ability to understand the correlation between players' actions and their private information. The two must be applied together in a situation where others' actions are not observed. If one of them fails, the individual will suffer the curse.

Let's take Groucho Marx as an example. Before applying, he could use contingent thinking to reason about the event of him being accepted as a member. Within this hypothetical scenario, he could infer from the acceptance that it is not a club he likes, which would lead him to resign. Thus, by applying both computations together he would conclude that it is optimal not to apply in the first place. This is not what he did, which allows us to conclude that he was unable to perform at least one of these two computations. On the other hand, after applying and being accepted he correctly infers that it is not a club he likes, and therefore resigns. It therefore seems he was able to perform inference from other's actions when these were not hypothetical. In this situation there were no unobserved states and therefore contingent thinking was not required.

The idea that players perform better in games where they can observe others' actions —i.e. in sequential games, as opposed to simultaneous games — seems intuitive and finds support in the literature. But there are different mechanisms related to cognitive computations that could have this effect on behavior. Eyster and Rabin (2005) [13] propose a theoretical model that explains suffering from the winner's curse *-cursedness-* purely as a failure of individuals making inferences from others' actions. 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 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 contingent thinking does not play a role and it is the differential ability of performing inference from others' actions in simultaneous vs. sequential situations that drives differences in behavior. In other words, inference from others' actions is switched off in simultaneous environments. 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. 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 curse. To do so we introduce the Barrel of Apples Game that can be played simultaneously, sequentially and in an extended version. We use a within-subject design and classify subjects in to behavioral types using their behavior in the simultaneous and sequential games. By correlating behavior in the extended version with subjects' behavioral types we can draw conclusions about what cognitive computations appear to drive behavior. We do so by manipulating one player's information structure and observing the effects on the other's decisions.

We find that about a third of our 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 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 are cursed in the sense of Eyster and Rabin, and changes in behavior cannot be explained by the mechanism they conjecture.

1.1 Literature Review

The intuitive idea that individuals perform better in sequential environments, compared to simultaneous ones, also finds empirical support in lab experiments. Kagel and Levin (2002) [20] summarize a series of studies in which they find that individuals perform closer to equilibrium predictions in ascending price compared to sealed-bid auctions. Ali et al. (2008) [4] include similar results when comparing simultaneous to sequential committee voting.

More recently, Esponda and Vespa (2014) [12] also find the same in a voting experiment specifically designed to study individual behavior in simultaneous vs. sequential play. They additionally show that sub-optimal behavior is robust to experience and even subjects receiving hints about pivotality. Their design provides a clean test to show that differences in performance are due to individuals' difficulties to perform the necessary cognitive calculations in simultaneous environments. Our design shares some features with EV, but it allows to distinguish whether these difficulties arise because subjects do not understand the correlation between the other's actions and private information, or whether they are related to cognitive thinking per se. We discuss the conceptual differences between the two papers in more depth in section 2.3. Differences in the design are discussed further in section 3.

Charness and Levin 2009 [9] use a version of the *acquiring a company game* (Bazerman and Samuelson 1983 [5]) and find that a significant portion of their subjects suffer the winner's curse. Using appropriate controls integrated in to their experimental design and alternative treatments they conclude that this can be attributed to failure to perform the appropriate contingent thinking. Their design, though, does not allow to compare subjects' performance in simultaneous versus sequential action environments. Furthermore, the design is such that there is no information structure driving an opponent's, so it is not possible to draw any conclusions regarding subjects' ability to make inferences from other actions. Ivanov et al 2010 [17] use the maximal game to show that it is unlikely for the winner's curse to be driven by beliefs.

We already mentioned Eyster and Rabin (2005) [13] whose concept of cursed equilibrium (CE) is based on the assumption that some individuals may not be able to perform inference from others' actions. Except for the quoted conjecture in their conclusions, they make no other distinction between simultaneous and sequential environments. Analogy based expectation equilibrium (Jehiel 2005 [18], Jehiel and Koessler 2008 [19]) is a closely related but more general concept.

A different family of models assumes a *Cognitive Hierarchy* (CH) of types (Camerer et al. 2004 [8]). An individual of level k best responds to the mistaken beliefs that others are distributed across all types of a level lower than k. In what is dubbed the *Level-k* model, individuals of level k believe all others are of level $k-1$ (Stahl and Wilson 1994 [24], 1995 [25], Nagel 1995 [22]). The idea is that some step-level reasoning can explain the winner's curse in some contexts (Gneezy 2005 [15], Crawfornd and Iriberri 2007 [11]) . Still, neither CH nor CE models can account for the change in behaviour that subjects display in our experiment, unless one assumes that one's level of sophistication is endogenous to the game played.

The idea of endogenous strategic sophistication has been explored experimentally by Agranov et al. 2012 [1] and Alaoui and Penta 2013a [2]. These experiments manipulate subjects' beliefs about their opponents and find that an individual's apparent level of sophistication is higher when she believes her opponent to be more sophisticated. Changes in strategic sophistication are not a result of changes in the game's format as in our case. Alaoui and Penta 2013b [3] propose a theoretical model in which individuals' depth of reasoning is a function of their cognitive abilities and payoffs. This model allows for the depth of reasoning to be affected by the cognitive demands of different games, but, according to their definition, the two game formats we use are "cognitively similar".

Finally, Brocas et al. 2014 [7] use mousetracking to uncover individuals' thinking process in private information betting games. They also detect heterogeneity in behavior that correlates with different modes of strategic thinking, which are identified through the mousetracking. However, the environemt of their design is not directly related to the common-value environmets in which the winner's curse is endemic.

2 The Barrel of Apples game

2.1 The game

For our experiment we use a novel game we call the "Barrel of Apples game" (BoA). Two players ($i \in \{1, 2\}$) are offered a barrel of 10 apples. The barrel may be good ($\theta = G$, contains 10 good apples) or bad ($\theta = B$, contains a number $Q \in \{1, \ldots, 9\}$ of bad apples). The prior probability for each of these events is $\frac{1}{2}$. Both players draw an apple from the barrel, check its quality and put it back. Notice that drawing a good apple ($s_i = \boldsymbol{g}_{ood}$) means it is more likely for the barrel to be good ($Pr\{\theta = 1 \,|\, s_i = g\} > \frac{1}{2}$). Drawing a bad apple ($s_i = \boldsymbol{b}_{ad}$) means the barrel is certainly bad ($Pr\{\theta = 1 \,|\, s_i = b\} = 0$).

Both players choose whether to accept or reject the barrel: $x_i \in X = \{\boldsymbol{A}_{ccept}, \boldsymbol{R}_{eject}\}$. 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.

The players' payoffs depend on the state and whether they obtain the barrrel. Players want to obtain a good barrel and avoid a bad barrel. Each player wins 100 points if he/she *obtains a good barrel* or *does not obtain a bad barrel*. In all other cases the payoff is 0. Table 1 shows players' payoffs analytically, conditional on the state of nature and their actions.

(a) The barrel is *Good*

		P 2	
		Accept	Reject
P 1	Accept	100 , 0	100 , 0
	Reject	0 , 100	0 , 0

(b) The barrel is *Bad*

		P 2	
		Accept	Reject
P 1	Accept	0 , 100	0 , 100
	Reject	100 , 0	100 , 100

Table 1: Players' payoffs conditional on the state and their actions. P 1 and P 2 refer to players 1 and 2 respectively.

In our experiment, the BoA game is played in three formats: as a *simultaneous* game, where both players decide without knowing the other's choice; as a *sequential* game, in which player 1 makes a choice first and player 2 can observe it before making his own; and in the *pay to observe* (PtO) format, where player 2, before observing his signal, has the choice of paying 10 points to observe player 1's choice. Notice that in the PtO format player 2 essentially has the choice of turning a simultaneous game in to a sequential one. In the sequential format, or after paying to observe player 1's choice, the game ends after an acceptance from player 1.[1] The simultaneous and sequential formats are used in the experiment to gauge subjects behavior and classify them in to the behavioral types we describe below. Behavior in the PtO format is then used to understand the cognitive computations these different types perform.

The barrel of apples game is designed to create a very simple winner's curse environment and is therefore purposely different from particular market or other institutions. Still, the strategic situation faced by players is essentially that of a posted-price common-value offer, and the strategic thinking required is similar to that in common-value auctions and jury voting.

2.2 Behavior in the simultaneous and sequential formats

2.2.1 Player 1

Given player 1's priority over player 2, her payoff depends only on the state of nature and her own actions. The best strategy for her is to follow her private information: if the apple drawn from the barrel is bad, reject the barrel; if the apple is good, accept. The format of the game is also irrelevant for player 1 since either way she cannot observe player 2's action.

Note that the optimality of this strategy is not dependent on risk preferences. After observing her private signal, player 1 faces a choice between two lotteries with the same possible payoffs — 100 or nothing — but different probability distributions. Choosing the lottery that assigns a higher probability on 100 is optimal for any type of risk preferences. It turns out that, by the same argument, risk preferences are not relevant for player 2's choices either.

A different potential source of bias could be the presence of other-regarding (social) preferences. For instance, if player 1 cared about player 2's payoff more than for her own she could reject the barrel after observing a good signal to benefit player 2. In our experiment we control for such behavioral factors by sustituting player 1 with a computer. The focus of the analysis is the behavior of player 2.

[1]In the discussion we sometimes consider the sequential game as ending only after a decision from player 2, which is non-consequential when it follows an acceptance by player 1. This maintains the symmetry between the two games and facilitates the presentation, without any effect on the analysis.

2.2.2 Player 2

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 reject the barrel independetly of his own private signal. This is true in both the simultaneous and the sequential game. This strategy does not guarantee player 2 winning the 100 points. If the barrel is good, then player 1 will accept it and player 2 wins 0 points. But in that case, player 2's choice will not have any effect on the final outcome.

This gives us a first theoretical benchmark for player 2's behavior. A strategically *sophisticated* individual in this role should reject the barrel in both formats of the game and after observing any possible private signal. [2]

The literature in experimental games teaches us to expect deviations from optimal behavior. Not all individuals posses the strategic sophistication required in games as this. Another usefull benchmark is then to think of the behavior expected by a *naïve* individual in the role of player 2 that ignores the strategic elements of the game. Disregarding the effect of player 1's choices on the final outcome, player 2 faces a simple decision problem as the one described for player 1. The natural strategy then is to follow his own private signal. Such a naïve player 2 acecpts after a good signal and rejects after a bad signal in both game formats.

None of these two benchmarks allows for behavior to be influenced by the game's format. In an informal sense, the simultaneous game appears more complex comprared to the sequential one. Therefore, intuitively one might expect behavior to be closer to optimal in the case of the "easy" sequential game and more naïve in the "hard" simultaneous game. This introduces a third benchmark for behavior. One in which player 2: behaves naïvely in the simultaneous game and simply follows his private signal; and behaves optimally in the sequential game, rejecting independently of the private signal he receives. In lack of a better term, and given the similarity of this behavior to that of the comedian in our introductory anecdote, we call this *Groucho Marx (GM)* behavior. Table 2 gives a summary of the three theoretical benchmarks for behavior that we use.

How do these benchmarks relate to existing models of behavior in games? So-

[2]One could perhaps imagine an individual rejecting irrespectively of his own private signal for reasons that are not related to the avoidance of the winner's curse. Labeling such behavior as sophistcated would in that case be misleading. On the other hand, we could not come up with a plausible argument (especially none backed by empirical evidence in the behavioral literature), of why an individual should behave like that for other reasons. Even so, it will become clear when analyzing the experimental data that such a potential mislabeling of behavior does not affect our results in any significant way.

Benchmark	Naïve				Sophisticated				GM			
	sim		*seq*		*sim*		*seq*		*sim*		*seq*	
	bad	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*	*bad*	*good*
frequency of reject %	**100**	**0**	**100**	**0**	**100**	**100**	**100**	**100**	**100**	**0**	**100**	**100**

Table 2: Theoretical benchmarks for player 2's behavior. Numbers indicate the frequency with which a player 2 of a specific type is expected to reject the barrel when playing the game several times, conditional on the private signal (*bad* vs. *good*) and the game's format (*sim*ultaneous vs. *seq*uential).

phisticated behavior essentialy corresponds to that predicted by *(Perfect) Bayesian Nash Equilibrium*. Non-equilibrium models that allow for heterogeneity in strategic thinking also predict this behavior for some types of individuals with higher levels of sophistication. Individuals of low sophistication levels in cognitive hierarchy and level-k models, and cursed individuals in the sense of Eyster and Rabin's "cursed equilibrium" all behave like our naïve benchmark.[3]. These models' higher sophistication types' behavior matches the sophisticated benchmark. None of these models predicts behavioral types corresponding to the intuitive behavior of our GM benchmark.

As we discuss in the next section, in the experiment we use a computer to simulate the behavior of player 1 and subjects know the rule used by the computer, which is to follow its signal. It can therefore be argued that these theories have little relevance for our experiment. This is because they assume that players fail to form correct beliefs, which in our experiment is not possible since they know the computer's behavioral rule. This argument is based on a very strict interpretation of these models' assumptions. In any case, in Louis (2012) [21] subjects play the simultaneous and sequential game against human opponents and exhibit the same types of behavior as in the current experiment. Our argument is slightly more involved. Even if these behavioral models are viewed as "as if" models that help organize experimental data, we present a very simple environment where they fail to do so. Furthermore, we help disentangle the specific cognitive failures that drive bahavior. This gives a clear direction in which behavioral game theorists should look for improved models.

[3]In the context of a level-k model and for our game, one could assume that level-0 types just follow their signal and thus match the naïve benchmark. Alternatively one could assume that these types play randomly and therefore it is level-1 types that follow their signal, since they are best responding to a randomizing player 1. Either way, some lower level type will match the naive benchmark. See Crawford and Iriberi 2010 [11] for a discussion of the appropriate level 0 assumption in common value auctions.

2.3 Cognitive calculations and the Pay to Observe format

Behavior in the first two formats of the game can be indicative of which cognitive calculations an individual performs. In particular we are interested in her ability to perform *contingent thinking* (CT) and *inference from others' actions* (IOA). Note that in the sequential game a player only needs to apply IOA. In the simultaneous game, CT and IOA are combined in order to calculate the optimal strategy.[4]

A sophistcated type seems able to perform both these calculations. Naïve types fail in IOA, given their behavior in the sequential game. Their behavior in the simultaneous game can then be attributed to a failure of IOA, perhaps compunded by a failure of CT. For GM types things are less clear: they perform IOA correctly in the sequential game, like sophisticated types, but fail in at least one of the two computations in the simultaneous game.

At this point the difference should be clear between CT and IOA and the two concepts in Esponda and Vespa (2014) [12] (EV): *hypothetical thinking* and *information extraction*. The former two refer to cognitive computations while the latter pair seem to refer mostly to the strategic situation in which the computations take place. Since in the sequential games (both the BoA game and the voting game in EV) the only relevant computation is IOA, it coincides with the concept of *information extraction*. On the other hand, the *hypothetical thinking* that takes place in the simultaneous games includes both CT and IOA. EV do not offer evidence on whether or not their subjects are performing IOA in the simultaneous games. The PtO format allows us to answer that in our design.

In the PtO format individuals take two descisions: first, whether or not to pay 10 points to observe player 1's choice; subsequently, they decide whether to accept or reject the barrel. Once faced with this second choice, individuals are essentialy in the same position as they are in the sequential or simultaneous game, depending on whether they paid or did not pay, respectively, in the first choice. Their behavior in this second decision is therefore expected to depend on their type in the same way as described above. How about their behavior in the first choice in this format?

A sophisticated type understands that it is optimal to reject the barrel independently on whether she actually observes player 1's choice or not. Paying to do so is therefore a waste. Naïve types do not behave optimally, but they only take in to account their

[4]Rangel et al. (2008) [23] propose a framework for the basic computations involved in making a choice. They brake down the process of value-based decision making in to five basic computations: *Representation*, *Valuation*, *Action selection*, *Outcome evaluation*, and *Learning*. The first four follow one another and the outcome of each one feeds in to the next, while Learning is based on the outcome of the process and feeds in to the other four computations when they are repeated in the future. Given our definitions, CT can be seen as part of the Representation process and IOA as part of Valuation.

own private information when deciding whether to accept the barrel. For them player 1's choice is irrelevant and therefore paying to observe it is useless. On the other hand, GM types behave differently in the simultaneous and the sequential game. Depending on what drives this difference in behavior, a GM type could consider observing player 1's choice as valuable.

There are two cases. In the first, a GM type fails in IOA when he cannot observe player 1's action. That is, he does not see the correlation between actions and player 1's private information. If that is the case, then changes in the informativeness of player 1's private information are not expected to have an effect on the GM type individual's willingness to pay to observe her choice. The second case is that the GM type can always perform IOA, but fails in CT. Through IOA he understands that player 1's information reveals her private information. The failure in CT does not allow him to see that his decision only matters when player 1 rejects. Therefore, he sees some value in actually observing player 1's action, and this value depends on the informativeness of player 1's signal. This informativeness is manipulated in the experiment through the parameter Q, the number of bad apples in the bad bag. This allows us to distinguish between the two cases.

3 Experimental Design

We use a within-subject design. In the experiment, subjects play variations of the BoA game repeatedly 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: accept when drawing a good apple, reject when drawing a bad apple. Subjects know this rule.

Subjects face all three game formats in the 40 rounds of play. 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 assure 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 5 decisions in the 20 sequential treatments.

The number Q 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 treatment. Table 3 shows all possible treatments and the number of times each subject is in one of them. The order with which subjects go through the treatments is randomised separately for each subject to control for order effects.

	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 3: Number of decisions per treatment for each subject

Apples are presented as balls in urns on the subjects' screens. They have different color 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 to avoid unintended connotations between colors and behavior. No feedback is given to subjects during the experiment.

Our design shares some features with the one of EV, but there are important differences. The main difference is of course the use of the BoA game, which has some advantages. First, it is a simpler game with a smaller number of possible states. Second, it is a game that could be played with human opponents that would have the same expected behavior as the one of the computer in our experiment.[5] This provides for a more natural strategic environment compared to one where the computer opponent behaves according to an arbitrary probabilistic rule. Such arbitrary rules could make subjects suspicious and introduce unwanted "experimenter demand" effects.[6] Finally, and more importantly, the BoA game can be played in the PtO format that allows us to answer the paper's main question, namely why individuals behave differently depending on the game format.

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 supplemental material. After reading the instructions and before starting the experiment, subjects had to pass a short comprehension test.[7] They were not allowed to continue until answering all questions correctly. Subjects had no significant difficulties in doing so.

To determine subjects' earnings, 7 of the 40 rounds were chosen randomly and the

[5]See Louis (2012) [21].
[6]See Zizzo (2010) [27]
[7]The five questions can be found in the supplemental material

subjects' outcome in these rounds was considered. An additional 100 points were given to each subject to assure there was no bankruptcy from 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.

The experiment was conducted in six sessions with a total of 97 subjects participating.[8] Subjects were recruited among undergraduate students at ETH Zurich and the University of Zurich using ORSEE (Greiner 2003 [16]). The experiment was conducted in the Experimental Economics Lab of the University of Zurich using Z-Tree (Fischbacher 2007 [14]).

4 Results and Analysis

Our analysis takes the following steps. We first present some data on aggregate behavior in all three fromats. We then use the data from behavior in the simultaneous and sequential formats to classify subjects in to different behavioral types. Next, we turn attention to behavior in the PtO format, broken down according to the classification obtained at the previous step. It is from this last step that we make our conclusions about the computations that drive subjects' behavior in the different formats. To keep the discussion more concise we now refer to the three game formats as *Sim*, *Seq* and *PtO* respectively.

4.1 Aggregate data

In *Sim* we always observe a decision to accept or reject. In *Seq*, when the computer accepts, there is no decision for the subject to take, hence no observation. In *PtO* we always observe a decision whether or not to pay for the option to observe the computer's choice. If the subject decides not to pay, we also observe his decision to accept or reject. As in *Seq*, if the subject pays we observe a decision to accept or reject only if the computer has rejected.

Table 4 shows the aggregate behavior of subjects across treatments. The frequency with which subjects reject the barrel is indicated, conditional on being in each format,

[8]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 to a failure of the recruitment system to send the automated reminder to subjects. Since the experiment consists in a series of individual decisions, with no interaction among subjects, the number of participants per session should not affect the results. Furthermore, the exact same instructions and procedures were used in each session. A pilot session conducted during the calibration of the design is not included in the analysis. The results we present are robust to omitting any particular session.

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.

On one hand, almost all subjects reject after observing a bad signal. This indicates that subjects understand that a bad signal unambiguously reveals the state to be bad. On the other hand, the frequency of rejection does not take an extreme value in any of the situations after observing a good signal. For the case of simultaneous decisions (*Sim* and *PtO* after not paying) subjects reject the barrel after a good signal in about a third of the cases. For the case of sequential decisions (*Seq* and *PtO* after paying) the frequency of rejecting after a good signal roughly doubles.

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

To establish the factors that determine whether subjects accept or reject the barrel we run probit regressions with *subject's choice* as a dependent variable. Results are presented in table 5. We pooled together all decisions from treatments in *Sim* and *Seq*. In *PtO*, *sequential* is endogenous, since subjects decide themselves whether to pay to observe the computer's choice, or not. We present the coefficients' estimates for *PtO* here for completeness, but these should be viewed with some reservation. We defer discussion of behaviour in *PtO* to subsection 4.3 and now turn our focus on subjects' behaviour in *Sim* and *Seq*.

The negative and strongly significant coefficient for *sequential* supports what is already evident in the aggregate data of table 4: subjects are more likely to play optimally and ignore their private signal after observing the computer reject. This is expected and is in line with previous results in the literature and our intuition.

The estimated coefficients of the other variables taken in to account here also conform to intuition. For *subject's signal* the coefficient is positive, strongly significant and has the highest magnitude. This reflects what is shown in table 4: in almost all cases, subjects

Probit regressions
Dependent variable: *subject's choice ('reject' = 0, 'accept' = 1)*

	Sim and Seq	PtO
subject's signal	2.326**	2.560**
	(0.179)	(0.246)
sequential	-0.758**	-0.581*
	(0.136)	(0.280)
Q	0.061**	0.082**
	(0.019)	(0.025)
round	-0.009*	-0.013*
	(0.004)	(0.006)
bad color	0.056	0.014
	(0.078)	(0.118)
Intercept	-1.972**	-2.369**
	(0.232)	(0.292)
# of obs.	1456	699

* p < .05 ** p < .01

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

Table 5: 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 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).

reject after receiving a bad signal. The remaining variables mostly capture variation when subjects receive good signals. Concerning Q, notice that as it increases private signals become more informative. The positive coefficient for Q indicates that subjects understand that and tend to follow their private signal more when Q is higher. Of course, while this seems natural, it is not optimal in this context.

One issue that was a concern given the high number of rounds was the possibility of subjects learning during the course of the experiment. This could lead them to change their behavior and make it harder to draw any inferences about what is driving it. There is no feedback about outcomes during the experiment, so any learning could only come through gaining experience with the game.[9] Such an effect is captured by the coefficient for *round*. While it turns out to be negative and weakly significant , it has a very small magnitude. In fact, keeping the other variables at their mean and increasing *round* by one standard deviation increases the probability of rejection by slightly over 3%. Furthermore, the coefficient is not significant when excluding the first or last 5 rounds of decisions. The other significant coefficients are robust to such a manipulation.

The aggregate results presented here hide a significant degree of heterogeneity in subjects' behavior. Each subject made 10 decisions in *Sim*, with an average of 7.6 made after receiving a good signal. Each subject made an average of 5 decisions in *Seq*, with an average of 2 made after receiving a good signal. For each individual we calculate the frequency of rejection in each situation, 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.

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 behaviours and some of the spikes in the distribution directly coincide with these points.

4.2 Classification

We classify subject's behaviour within and across game formats according to three types, corresponding to the three behavioral benchmarks, and an additional random type added for robustness. The fact that the distribution of behaviour shows spikes in the points corresponding to these prototypes indicates that this choice may not be completely misguided. Still, most of the data does not lie on any of these points so our

[9] See Weber 2004 [26].

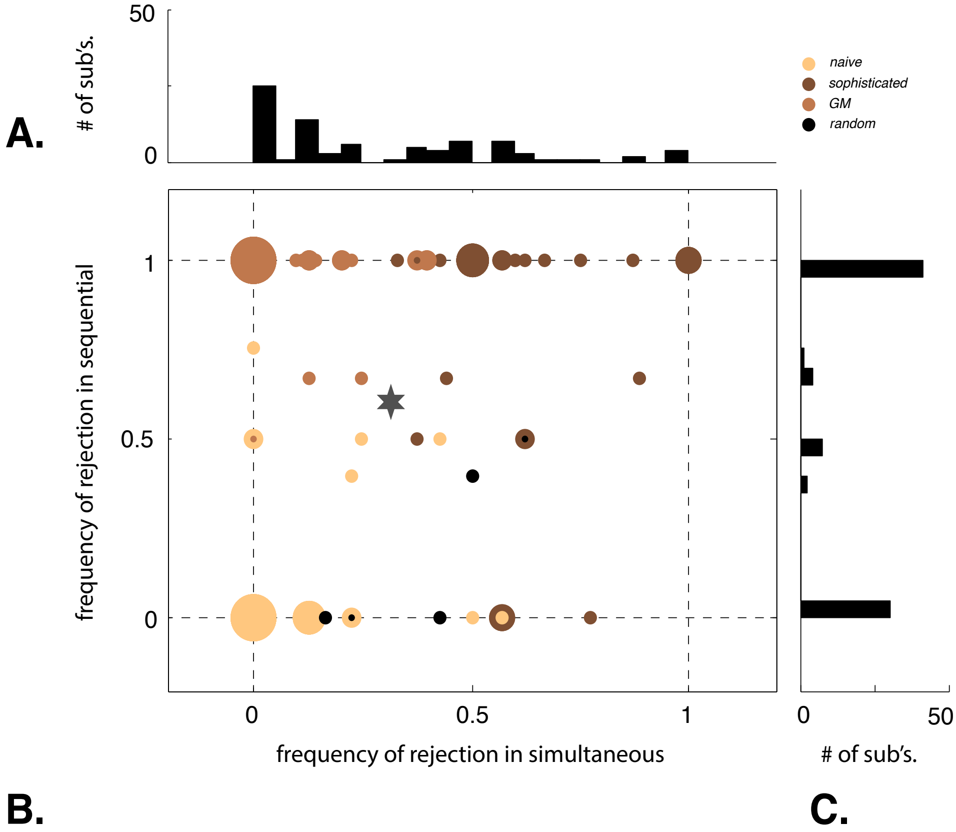

Figure 1: Individual frequency of rejection in situations 1 and 2. Each circle is centred at a point with coordinates corresponding to the frequency of rejection in the two situations. The radius of the circle is proportional to the number of subjects whose frequencies coincide. Colors refer to the types 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.

classification scheme must deal with the noise in the data.

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.[10]. It would therefore not be possible to classify these individuals using this method.

The method we use instead is distribution-based clustering. Choice distributions are derived from a hierarchy of logit choice models.[11] Such structural models are often used to organize experimental data in the literature regarding cognitive hierarchy and level-k models.[12] We use such a model here because the predicted behavior for different types matches that of the behavioral types we posit, allowing for noise.

We 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 noise parameter μ, which enters the logit response function. We then use an expectation-maximisation type algorithm to find the value of μ that maximizes the likelihood that the data is explained by our model. At the same time the algorithm classifies subjects into different types.[13] A formal specification of the model is given in the supplemental material.

Naïve	**Sophisticated**	**GM**	**Random**	μ	(st. error)
30	32	32	5	14.979	(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.

Table 6 summarizes the results of the classification excersise. The vast majority of 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.[14]

[10]These observations are not included in the scatter plot.
[11]See Stahl and Wilson (1994) [24].
[12]See Crawford et al (2010) [10] and references therein.
[13]See Bhatt et al. 2010 [6] for a similar application.
[14]There are cases of subjects that have very similar, or even the exact same, frequencies of rejection in the

The exact specification of the model could be modified slightly — e.g. different beliefs for sophistcated types or not allowing for random types. This would not affect results significantly. The specification used gives the best fit to the data among these alternatives. Figure 2 compares the data to the predicted frequency of rejection for the four possible format - signal combinations. In general the model does a good job in fitting the relative frequencies in each of the four cases. The model over predicts by a small degree the frequency of rejection in the simultaneous game after receiving a good signal. Other than that, the fit is good.

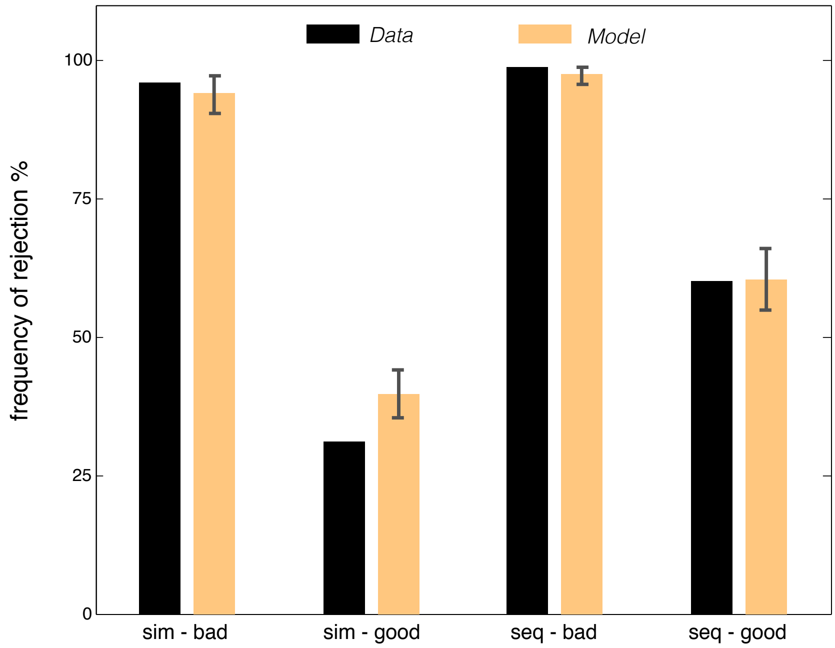

Figure 2: Model fit to the data. Error bars give 95% confidence intervals. These are calculated using Monte Carlo simulations, as explained in the supplemental material.

4.3 Behavior and computations

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.

two formats but are classified as different types. This happens because the classification algorithm takes in to 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.

	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 4 that subjects pay for this option on average 37.7% of the time. Table 7 breaks down this average for each type.[15] While naïve types pay to observe slightly more often than sophistcated types, the average for both is much lower than the one for GM types. 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. 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.

We are now in a position to answer the paper's main question. What drives the differences in behavior between the simultaneous and sequential formats? To begin with, it seems that the differences observed in aggregate behavior are created by a particular subset of subjects, the ones we classify as GM types. So the question is answered by undestanding why these subjects behave differently: are they able to perform IOA in *Seq*, but lose this ability in *Sim*? or is it the case that they can always perform IOA and their sub-optimal play in *Sim* is the result of a failure to perform CT? If the former is true, then GM types do not realize the information contained in player 1's choice before actually observing it. Their willingness to pay to do so should not be affected by changes in Q, that affect the informational content of said choice. The opposite should hold if the latter is true.

The evidence from our experiment supports the second posibility. Figure 3 presents

[15]We do not include the 5 subjects classified as random types in the remaining analysis.

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.

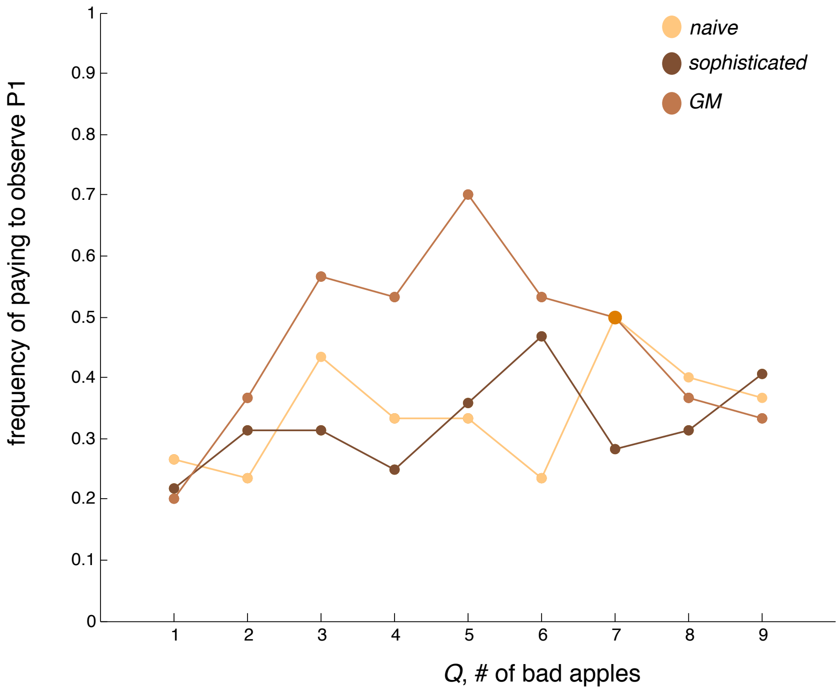

Figure 3: Paying to observe player 1's choice, by type and as a function of Q.

This pattern of behavior 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 means that GM types believe that by paying to observe the computer's choice is equivalent to paying to observe a *second private signal*, additional to their own. Their willingness to pay therefore depends on the value of this additional signal. This value depends on Q. For a single private signal, the higher Q, the higher its value. In particular, Making a decision after observing a private signal, compared to doing so with no private information, results in a positive difference of $100 \times \frac{1}{2} \times \frac{Q}{10}$, which is linear and increasing in Q. For an additional private signal, which is the relevant case here, the value as a function of Q is first increasing and then decreasing. Formally, the difference in expected value between a decision taken after observing two signals compared to one taken after observing only one is $100 \times \frac{1}{2} \times \frac{Q(10-Q)}{10}$. This has an inverted-U shape with a global maximum at $Q = 5$.

We 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 3: 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 $Q = 5$.

It turns out 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 may seem that subjects of these types eventually understand that it is better not to pay. Although, if that is true, it is not clear why the same subjects fail to correct their behavior to the same extend during the experiment when facing the choice of accepting or rejecting the barrel. We think that a different explanation is more plausible. 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 the subject will most likely not be given the chance to accept or reject the barrel, which can be frustrating. It seems more likely that these subjects choose to pay to observe less frequently in order to avoid this frustration.

5 Conclusions

This paper adds to the effort to understand individuals' susceptibility to the winner's curse, and how this is driven by the failure of specific cognitive computations: contingent thinking and inferences from others' actions. To this end we design and run a lab experiment using a simple game that incorporates a winner's curse environment. We find substantial heterogeneity in how subjects in our experiment react to differences in the game's format. One subset does not have substantial difficulties 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 the result of a failure to perform inference from others' actions, possibly compounded by a failure of contingent thinking. The difference in aggregate behavior across game formats is almost entirely attributed to a third subset of subjects.

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).

Our design leads us to conclude that these individuals have no problem in performing inference from others' actions, neither in the simultaneous or the sequential game formats. They are not "cursed" in the sense of Eyster and Rabin (2005). Their behavior is 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 increasing popularity of models of strategic thinking that incorporate such heterogeneity. 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 behavioral 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, we show evidence that the underlying differences in cognitive abilities 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

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

[2] Larbi Alaoui and Antonio Penta. Depth of Reasoning and Incentives. unpublished, 2013.

[3] Larbi Alaoui and Antonio Penta. Endogenous Depth of Reasoning. unpublished, 2013.

[4] 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.

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

[6] 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.

[7] Isabelle Brocas, Juan D. Carrillo, Stephanie W. Wang, and Colin F. Camerer. Imperfect Choice or Imperfect Attention? Understanding Strategic Thinking in Private Information Games. *The Review of Economic Studies*, 2014.

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

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

[10] 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.

[11] Vincent P. Crawford and Nagore Iriberri. Level-k Auctions: Can a Nonequilib-

rium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in PrivateValue Auctions? *Econometrica*, 2007.

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

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

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

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

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

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

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

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

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

[21] Philippos Louis. Information, Behavior, and the Design of Institutions. *Doctoral dissertation*, 2012.

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

[23] Antonio Rangel, Colin Camerer, and P. Read Montague. A framework for studying the neurobiology of value-based decision making. *Nature Reviews: Neuroscience*, 2008.

[24] Dale O. Stahl and Paul W. Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 1994.

[25] Dale O. Stahl and Paul W. Wilson. On Players Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 1995.

[26] Roberto A. Weber. 'Learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 2003.

[27] Daniel John Zizzo. Experimenter demand effects in economic experiments. *Experimental Economics*, 13(1):75–98, 2010.