Trading Participation Rights to the Red Hat Puzzle. Will Markets allocate the rights for performing decision tasks to the more abled players?

Choo, Lawrence C.Y

University of Exeter

27 April 2014

Trading Participation Rights to the Red Hat Puzzle. Will Markets allocate the rights for performing decision tasks to the more abled players? *

Lawrence C.Y Choo †

This Version: April 2012

Abstract

This paper investigates the conventional wisdom that markets would naturally allocate the rights for performing decisional task to those players who might be best suited to perform the task. We embedded the decisional tasks in a stylised setting of a game, motivated by Littlewood (1953) Red Hat Puzzle, when the optimal choices in the game require players to employ logical and epistemological reasoning. We present a treatment where players are permitted to trade their participation rights to the game. The payoffs are furthermore calibrated such that the players who know the optimal choice in the game should value the rights strictly more than those who do not. However, aggregated performances in this treatment were found to be significantly lower than the control treatments where players were not permitted to trade their participation rights, providing little support for the conventional wisdom. We show that this finding could be attributed to price "bubbles" in the markets for participations rights.

Keywords: Game Theory, Experimental Economics, Financial Markets

JEL Classification: C72, C92, G34

*We are grateful to Todd R. Kaplan, Shyam Sunder, Peter Bossaerts, Elena Asparouhova, the conference participants of Experimental Finance Tilburg 2013 and ESA World Meetings 2013, for their valuable suggestions and comments. We acknowledge the financial support from the University of Exeter Business School.

†Department of Economics, University of Exeter, Rennes Drive, Exeter EX4 4PU, United Kingdom. *cylc201@exeter.ac.uk*

1 Introduction

Most societies integrate some form of design where individuals are able to buy, and sometimes even sell the “rights” for performing specific decisional tasks. An early example from the 17th to 19th centuries, is the British Army’s purchase system, where commissioned ranks and responsibilities were sold at pre-determined prices (Bruce, 1980; Brereton, 1986). Another example that should be more familiar in this present era, is the market for corporate governance, where managers compete for the rights to manage the corporate resources of a targeted firm (Jensen and Ruback, 1983). The conventional wisdom in the above examples is the idea that when properly structured, markets should allocate the rights for performing each decisional task to those individuals might be best abled to perform the task. When payoffs in each task are strictly increasing with performance, the economic intuition here is simply that individuals who are better abled to perform the task would value the rights more than individuals who are less abled to perform the task. Given the differences in valuation, markets should therefore facilitate the transfer of rights to those who value them the most. However, the concern for any economic designer such as a social planner, regulator or manager, is whether the conventional wisdom *would* naturally hold in most circumstances. More specifically, would markets naturally result in the allocation of rights to the more-abled individuals? We believe that some resolution to this question would provide valuable insights for the designer. This paper thus presents an experimental design to put the conventional wisdom to the test.

We embed the decisional task for individuals in the stylised setting of a game, motivated by the Red Hat puzzle (RHP), a logical reasoning problem attributed to Littlewood (1953). In our TRADE treatment, individuals are presented the opportunity to trade their participation rights to the game. To do so, we first endow each individual with one “token” and some information about their decisional task in the game. Here, each token represent an individual’s participation rights to performing the task in the game. Thereafter, individuals enter a market where they trade tokens but only amongst the other individuals with the same information - individuals will be “competing” for identical decisional tasks. After all trades are completed, individuals who had sold their token are compensated by the sales revenue for giving up such rights and only individuals with *at least one token* will proceed into the game. At the end of the game, individuals’ tokens will be redeemed at a rate that depends on the outcomes of their choices in their decisional task within the game.

Thus purchasing another individual's token not only buys over his participation rights in the game, but also his potential payoffs in the game.

The RHP has many attractive features for the purpose of our study, especially as to how individuals' abilities would influence their payoffs

(i) The equilibrium analysis of the game leads to a unique set of optimal choices in each decisional task, that will result in that individuals' tokens being redeemed at the Pareto optimal rate of β^*. Thus individuals' performances becomes a binary outcome as to whether he had adhered to the optimal choices or deviated.

(ii) The optimal choices are non-trivial nor obvious and individual's *sophistication* (e.g., strategic thinking, cognitive reasoning abilities, intelligence, problem solving skills) are integral in knowing the optimal choices.

(iii) The complexity of the decisional task for individuals in the game depend on the states of nature in the game, providing variations in the study of individual choices.

(iv) The optimal choices are independent of whether individuals are first permitted to trade tokens or the number of tokens owned, allowing for direct comparisons of TRADE with the control treatments where individuals are each endowed with one token but are prohibited from trading tokens.

If the equilibrium analysis suggest there to be no theoretical difference between optimal choices in TRADE and the control treatments, how would the experimental design test the conventional wisdom put forth in the beginning of this section?

Weber (2001, Experiment 2) and Bayer and Chan (2007) used the RHP to study level-k (Nagel, 1995; Stahl and Wilson, 1994, 1995; Costa-Gomes et al., 2001) reasoning behaviours in experimental research. Weber's research focused on the aggregated rate of adherence to the optimal choices. In the most trivial decisional task for subjects in the RHP, the adherence rates was found to be unity. However, the adherence rates was observed to decrease significantly as the complexity of the decisional tasks increases.[1,2] Although subjects choices in our experimental design may involve elements of level-k reasoning, we will omit such discussions as they divert attention from the main area of interest in this paper. Nevertheless, Weber's experiments point to heterogeneity in subjects' sophistications with respect to their choices in the RHP.

In a separate area of research, Kluger and Wyatts (2004) presented an innovative experimental design to study how heterogeneity in individuals' sophistication might affect market prices. To do so, they embedded the decisional task of the Monty Hall problem into an asset market experiment.[3] Their design could be summarised with the following thought experiment. Assume that there exist an asset that allows you to switch doors in the Monty Hall problem for a winning prize of \$100 - after you had made your initial choice and the non-prize door is opened. A unsophisticated individual would wrongly judge the probability of winning the prize through switching door at 0.50 and value the asset at \$50. A sophisticated individual would realise that the probability of winning the prize through switching doors is in fact 0.67 and value the asset at \$67. Focusing on mean prices, Kluger and Wyatts (2004) findings suggest that when all subjects in the market (6 subjects each market) were unsophisticated - as judge by their behaviours in the Monty Hall problem, the mean price in the market was close to 50. However, when there was at least two sophisticated subjects in the market, the mean price was close to 67. Their research however did not focus on the allocation of assets but suggest that individuals' sophistication could be strong determinants of their pricing behaviours.

Building on these previous research, we are now in the position to present a behavioural prediction based on the allocation properties of markets, as to how the outcomes in TRADE might differ from those in our control treatment. Given that adherence to the optimal choices corresponds to

[1] Weber (2001) results could be of independent interest as his subject pool included Caltech undergraduate and graduate students. Caltech students are often known for their skills in logical reasoning problems (Camerer, 2003).

[2] Bayer and Chan (2007) results are slightly more difficult to interpret as they reported on "rationalizable behaviours". Such behaviours might also include those that are inconsistent with the predicted behaviours.

[3] The Monty Hall problem is from the TV gameshow "Let's Make A Deal" where the Host, Monty Hall, hides a winning prize behind three closed doors. A contestant is invited to choose one of the doors to open, but before doing so, Monty is committed to opening a non-prize door. Thereafter Monty presents the contestant the opportunity to switch their choice to the other unopened door. The dominant strategy here is for the contestant to always switch since the probability of winning the prize by doing so is 2/3.

the tokens being redeemed at the Pareto optimal rate β^*, the sophisticated individuals (those who have sufficient abilities to know the optimal choices in the game) should therefore value the tokens strictly more than the unsophisticated individuals (those who have insufficient abilities to know the optimal choices in the game). This difference in valuations leads to a behavioural prediction that when presented the opportunity to trade tokens, such as in TRADE, it becomes incentive compatible for sophisticated individuals to purchase tokens and unsophisticated individuals to sell tokens.[4] When this occurs only sophisticated individuals should therefore enter the game.

To make comparisons between the performances of across treatments, we will focus on the *efficiency rate* - the ratio of tokens redeemed at the equilibrium rate β^*. If markets do result in the allocation of tokens to the sophisticated individuals as suggested by our behavioural prediction, we should therefore expect the efficiency rate in TRADE to be significantly higher than those in the control treatments. Thus comparisons between treatments should provide some insights to the conventional wisdom put forth at the beginning of this paper.

Our experimental results can be summarised as followed. The efficiency rate in TRADE was found to be significantly lower than the control treatments, providing little support for our behaviour prediction. To provide some explanation to this finding, we first studied the mean prices of tokens in the markets of TRADE. Here, mean prices were often observed to be above β^*. Such price "bubbles" may have important implications on the allocative outcomes of markets as at any price above β^*, sophisticated individuals should strictly prefer to sell their tokens and avoid the decisional task in the game altogether. This is consistent with our subsequent observation, where subjects in TRADE who had purchased additional tokens where proportionally less frequently found to have adhered to the optimal choices. However, we are also able to show that individuals' sophistication can be strong determinants of their pricing behaviours. Subjects in TRADE who had on average purchased tokens at prices depicted by our behavioural prediction, were observed to be proportionally more frequently found to have adhered to the optimal choices. Some support for these observations were provided in our econometric analysis. Thus, our data suggest that markets do not naturally result in the allocation of rights to the more abled individuals.

[4] The difference in sophistication side-steps Milgrom and Stokey (1982) no trade theorem, as both the sophisticated and unsophisticated individuals can have expected gains from trade.

The rest of this paper is organised as followed, individuals will be referred to as "players". Section 2 introduces our experimental design, Section 3 presents the equilibrium analysis of our treatments, Section 4 presents our behavioural prediction, Section 5 formalises our hypotheses test, Section 6 describes our experimental procedures, Section 7 presents our experimental results and finally, Section 8 concludes.

2 Experiment Design

Three treatments are considered in this paper, BASE1, BASE2 and TRADE. However, only in TRADE were players allowed to trade tokens. To motivate our experimental design, we shall first present an overview of the RHP. Thereafter, we will introduce a generalised framework that is applicable to all treatments. Finally, given this generalised framework, we will show how the treatments vary.[5]

2.1 The Red Hat Puzzle

The RHP and its variations are commonly found is most graduate level game theory textbooks (e.g., Fudenberg and Tirole, 1993; Myerson, 1997; Maschler et al., 2013), discussions about common knowledge (Geanakoplos and Polemarchakis, 1982; Geanakoplos, 1994) and epistemological reasoning (Fagin et al., 1995). It is often described with three girls, each wearing a coloured hat - red or black, seated around a circle. Each girls sees all other hats but her own - all hats are black. An observer remarks that "there is at least one black hat" and asked the first girl if she knew the colour of her hat, to which she replied (publicly) with "No". The observer asked the second girl whom again replied with "No". However, when the observer asked the third girl, she replies with "Black". How did the third girl know her hat colour?

Let us first consider the case where only the first girl was wearing a black hat. Here, the first girl would immediately reply with "Black" since she does not see any other black hats. Observing the reply of the first girl, the second girl reasons that the first girl must have observed no other black

[5]Due to the treatments considered in this paper, the experimental design involves features that are different from those previous adaptions of the RHP by Weber (2001) and Bayer and Chan (2007). Thus direct comparisons to their results will not be prudent.

hats, and deduces her hat to be red - she replies with "Red". The same logic applies to the third girl. Now consider the scenario where the first and second girls were wearing black hats. The first girl remains uncertain and replies with "No". The second girl reasons, that the first girl must have seen another black hat and therefore deduces her own hat to be black - she replies with "Black". The third girl reasons, that the second girl must have only observed one other black hat (the first girl's hat) and hence deduces here own hat to be red. She thus replies with "Red". Now returning to the initial illustration, the third girl observed that the second girl had replied with "No". She therefore reasons, that the second girl must have seen two black hats, and therefore deduces her own hat to be black.

Each girl in the RHP faces the decisional task of ascertaining her own hat colour, and she does so through a process of logical and epistemological reasoning.[6] Furthermore, the task becomes more complex and challenging as the number of black hats observed increases. For these reasons, sophistication in integral in accomplish the task within the RHP. The challenge in this paper is to modify the RHP into a design which allows us to incorporate treatment variations, whilst still retaining the main characteristics of the game. For this we refer to the generalised framework.

2.2 Generalised Framework

The generalise framework will consist of two distinct stages, the pre-game stage, where players trade tokens, followed by the game stage, where players perform their decisional task in the settings of a game. Let $1_G \in \{0, 1\}$ be an exogenous parameter that determines if players are permitted ($1_G = 1$) to enter the pre-game stage. This implies that if $1_G = 0$, players go directly into the game stage and if $1_G = 1$, players' participation in the game stage will depend on their decisions in the pre-game stage. The generalised framework will begin with following parameters.

There are $N = \{1, 2, ..., n\}$ set hats with $M = \{1, 2, ..., m\}$ set members in under each hat. Let player i_j refer to the $j \in M$ member of hat $i \in N$. Nature chooses the true state $s \in S \equiv \times_{i \in N} H_i \setminus \{R_1, R_2, ..., R_n\}$, where $H_i \in \{B_i, R_i\}$ denotes hat i's colour - Black(B) and Red(R). There exist a common prior over S where each state $s' \in S$ is equally likely. For any state $s \in S$,

[6] Geanakoplos (1994) describe such a process as one of *indirect communication*, where each girl through their replies, communicate some information on the posterior about the true state of nature.

denote $Y(s) = \{1, 2, .., y\} \subseteq N$ as the set of B hats.

Each player observes all other hats' colour but his own. Denote $b_{i_j}(s) \in \{0, 1, 2, ..., n-1\}$ as the total number of B hats that player i_j observes for any $s \in S$ - this refers to player i_j's private information.[7] In addition, they are also publicly informed that the true state consist of *"at least one B hat"*. Since players under the same hat must make the same observations, $b_{i_j}(s) = b_{i_{j'}}(s) = b_i(s)$ for all $j, j' \in M$, $i \in N$ and $s \in S$. Finally, each player is endowed with one *token* and a working capital of $\bar{L} \gg 0$, issued as an interest-free loan.

2.2.1 The Pre-Game Stage

The pre-game stage consist of n markets in simultaneous operations, where players trade tokens but only with the other players under the same hat. In the absence of short-sales, let $p_i \geq 0$ denote the token transaction price in market $i \in N$, $x_{i_j} \in \{0, 1, 2, ..., m\}$ denote player i_j's after transaction inventory of tokens and $L_{i_j} \geq 0$ denote player i_j's after transaction holding of capital. Assume that token inventories are public information and $\bar{L}$ is sufficiently large to never be binding. If $x_{ij} = 0$, players' payoffs are immediately computed - to be discussed later.

2.2.2 The Game Stage

Only players with $x_{ij} > 0$ tokens may enter the game stage, where the face each face the decisional task of resolving their hat colour. There are $t = 1, 2, .., n+1$ discrete periods, where at each period $t < n+1$ players are simultaneously presented with the question *"Do you know your hat colour?"* to which they must independently and simultaneously reply with the following actions: "My Hat is R" (a_r), "My Hat is B" (a_b) or "No, I don't Yet Know" (a_n). The rules are such that each player (and that player only) ends the game stage at the period t_{i_j} whereby the actions $e_{i_j} \in \{a_r, a_b\}$ were chosen. This implies that players only proceed to the next period if he had chosen a_n in the previous period. To ensure that all players must eventually leave the game stage, players can only choose from the actions a_b and a_r if they make it to the $n+1$ period. Finally any action chosen in

[7]Alternatively, one could employ Aumann (1999) *sematic* approach where each player's knowledge of the true state is represented by the information partition $\mathcal{P}_{i_j}$ over S. Such an approach might be more precise but it makes the discussion more taxing with no obvious benefits. Nevertheless the analysis will be identical.

period t will only be public information in period $t+1$.

2.2.3 Payoffs

Players' payoffs (Π_{i_j}) are computed when they have either ended the pre-game stage with $x_{i_j}=0$ tokens or ended the game stage with actions $e_{i_j} \in \{a_r, a_b\}$. Here, the true state of nature is revealed, the players' loan ($\bar{L}$) are repaid and their tokens are each redeemed at the heterogenous rate $\beta(\mu, \delta, \alpha, H_i, t_{i_j}, e_{i_j}) \geq 0$ - in a slight abuse of notation we will write $\beta(\mu, \delta, \alpha, H_i, t_{i_j}, e_{i_j})$ as β_{i_j}. Table 1 depicts the generic tokens redemption rate for each player i_j, where $\mu > (1/2)\alpha > (n+1)\delta > 0$. The redemption rate can be summarised as followed: Each token has an initial value of μ that decreases by δ each time the player chooses a_n. In addition, the token's value decreases by α if he had incorrectly guessed his hat colour - choosing a_b (a_r) if $H_i = R_i$ ($H_i = B_i$). The payoffs are therefore determined as followed[8,9]

$$\Pi_{i_j} = \begin{cases} (L_{i_j} - \bar{L}) + \beta_{i_j} x_{i_j} = \beta_{i_j} & \text{if } 1_G = 0 \;\&\; x_{i_j} = 1 \\ (L_{i_j} - \bar{L}) + \beta_{i_j} x_{i_j} = p_i + (\beta_{i_j} - p_i) x_{i_j} & \text{if } 1_G = 1 \;\&\; x_{i_j} > 0 \\ (L_{i_j} - \bar{L}) = p_i & \text{if } 1_G = 1 \;\&\; x_{i_j} = 0 \end{cases} \tag{1}$$

Players who sold their tokens are thus compensated by the sales revenue of p_i for avoiding the game stage. Players who purchased additional tokens, not only buy over the other players participation rights to the game stage, but also their potential payoffs in the game stage. This is simply due to the fact that players' payoffs are dependent on the number of tokens owned and the redemption rate for each token. This completes the description of the generalised framework.

2.3 How the Treatments Vary

For any fixed $n \geq 2$, variations in the generalise framework can be achieved by specifying the number of members under each hat (m) and whether players are permitted to enter the pre-game

[8] When $1_G = 0$, we must have it that $\bar{L} = L_{i_j}$ and $x_{i_j} = 1$ since players are not permitted to enter the pre-game stage.

[9] Since players are each endowed with one token, their net transactions in the pre-game stage can be denoted as $v_{i_j} = x_{i_j} - 1$, where the market clearing conditions require that $\sum_j v_{i_j} = 0$ for all $i \in N$. As such, we can rewrite players' holding of capital as $L_{i_j} = \bar{L} - p_i v_{i_j}$.

Table 1: Generic Token Redemption Rate (β_{i_j}) for Each Player i_j

	$H_i = B_i$	$H_i = R_i$
$e_{i_j} = a_b$	$\mu - \delta(t_{i_j} - 1)$	$\mu - \delta(t_{i_j} - 1) - \alpha$
$e_{i_j} = a_r$	$\mu - \delta(t_{i_j} - 1) - \alpha$	$\mu - \delta(t_{i_j} - 1)$

stage (1_G). The three treatments are differentiated as followed:

BASE1: $n = 3$, $m = 1$ and $1_G = 0$.

BASE2: $n = 3$, $m = 6$ and $1_G = 0$.

TRADE: $n = 3$, $m = 6$ and $1_G = 1$.

Players in BASE1 and BASE2 hence always enter the game stage with exactly one token. BASE1 refers to the primitive description of the Red Hat puzzle. TRADE is the central interest of this paper, where players are permitted to trade tokens in the pre-game stage. Since TRADE and BASE1 differ on both m and 1_G, BASE2 was introduced to control for any potential difference that were driven by changes in m.

3 Equilibrium Analysis

The equilibrium analysis assumes that all players are "Rational" and "Sophisticated", and this being a common knowledge fact. Adapting the description put forth by Myerson (1997), we refer to rational players as those who seek to maximise their own payoffs. Similarly, we refer to sophisticated players as those who knows everything there is to know about the game and makes the same logically conclusion or consequences as a designer of the game would make. Therefore, although players may start the treatment uncertain of their hats' colours, they are assumed to always know the process of ascertaining their hats' colours in the game stage. Furthermore, players are also assumed to be risk-neutral. As it will be clearly in the relevant subsections, the equilibrium analysis of the respective treatments will show the following:

(i) *The optimal choices for any player in the game stage will only depend on $b_i(s)$ and will be independent on whether the players are allowed to enter the pre-game stage, the number of* *members (m) under each hat and the number of tokens own (x_{i_j}) by players in the game stage.*

(ii) *The equilibrium payoff for any player is $\Pi^*_{i_j} = \mu - b_i(s)\delta$ and will be independent of the treatment variations and whether players had entered the game stage.*

We will first detail the equilibrium analysis in BASE1 and thereafter extend the discussions to BASE2 and TRADE. Finally, we will show the equilibrium payoffs for any player in either treatments.

3.1 Equilibrium Analysis in BASE1

Players enter the game stage with exactly one token as since by definition, $\Pi_{i_j} = \beta_{ij}$, players should seek to maximise their token redemption rate. The equilibrium predictions here are for players observing $b_i(s)$ to ascertain their hats' colour at period $t^*_{i_j} = b_i(s) + 1$. The corresponding optimal choices are for players to choose a_n at all periods $t < t^*_{i_j}$, and at period $t^*_{i_j}$, choose a_b if $i \in Y(s)$ and a_r if $i \notin Y(s)$. Adherence to the optimal choices will result in players' tokens being redeemed at the rate $\beta^*_{i_j} = \mu - b_i(s)\delta$. Furthermore, adherence to the optimal choices will be Pareto optimal for all players.

To see why this might be so, let us first consider the dominant action for players at any period t, where they are certain or uncertain of their hats' colour. In the case where players are certain, the dominant action is obvious. Given the token redemption structure, they should choose a_b or a_r if they know their hats to be B or R respectively - choosing a_n incurs an additional cost of δ with no obvious benefits. What is less obvious is the dominant action at any period t for uncertain players. By Bayes rule, uncertain players must hold equal posterior to being under either hat colours - this will clear in the later discussions. Here, players face inter-period tradeoff between (OptionA) Ending the game stage at period with $e_{i_j} \in \{a_b, a_r\}$ or (OptionB) Choosing a_n and ascertaining their hats' colour at some later period $t' = t+1, t+2, .., n+1$. The expected token redemption rate for pursing OptionA will be $\mu - \delta(t-1) - (1/2)\alpha$. Since players are assume to always know the process of ascertaining their hats' colour, the expected token redemption rate for pursing OptionB will be $\mu - \delta(t'-1)$. Since by definition, $(1/2)\alpha > (n+1)\delta$, the expected token redemption rate for any $t' = t+1, t+2, .., n+1$ with OptionB will be strictly greater than those of OptionA. Thus uncertain players would choose a_n.

Having established the dominant action at any period for certain and uncertain players, we are now in the position to described the process by which players ascertain their own hats' colour. To show this process, we will return to the example introduced in section 2 of this paper, where $n = 3$ and $s = B_1B_2B_3$.

Each player begins period 1 of the game stage observing $b_i(s) = 2$ and remains uncertain. Since each state in S is equally likely, by Bayes rule, their conditional posterior of being under a B hat must be 1/2. Given the public announcement, it can only be common knowledge that there is *at least one B hat.*[10] However, each player privately knows there to be at least two B hats. Uncertain players in period 1 thus choose a_n. At period 2, having observed the previous period's action of each other players, each player reasons that if there was only one B hat, then some player must have observed no B hats, ascertained his hat's colour to be B and choose a_b in period 1. Since no one had done so, there cannot be only one B hat in the true state. Of course each player already knew this and there should be no revisions to their posteriors. Again uncertain players choose a_n. Finally at period 3, each player reasons that if there were only two B hats, then some players must have observed one other B hats, ascertained their hats' colour to be B, and choose a_b in period 2. Since no player had done done so, there cannot be only two B hat in the true state. Furthermore, since $b_i(s) = 2$, each player deduces their hats to be B and thus choose a_b. Players leave the game stage in period 3 and their tokens are each redeemed at the Pareto optimal rate $\beta^*_{i_j} = \mu - 2\delta$.

The above illustration can be extended to any $n \geq 2$ hats and the process will be similar. The astute reader should note the role that the common knowledge conditions play in the equilibrium analysis. In the absences of the common knowledge conditions, players cannot exclude the possibility that an action chosen by some other player is due to unsophisticated motivations or irrational behaviours.

[10] Alternatively, Aumann (1976) agreement theorem, show that the only event in S which can be commonly knowledge must include the entire states of nature S.

3.2 Equilibrium Analysis in BASE2

Players in BASE2 always enter the game stage with one token and only differs from BASE1 in the number of members under each hat. However, players in under each hat have the same private information, face the same decisional task and choose their actions both independently and simultaneously. This implies that the optimal choices for each player must be identical at all periods, for players under the same hat - choose a_n at all periods $t < t^*_{i_j}$, and at period $t^*_{i_j}$, choose a_b if $i \in Y(s)$ and a_r if $i \notin Y(s)$. Thus, increasing the number of members under each hat, has no implications on the optimal choices in the game stage and adherence will result in tokens being redeemed at the rate $\beta^*_{i_j} = \mu - b_i(s)\delta$.

Some Comments about the Equilibrium Analysis in BASE2

Since the optimal choices for players under the same hat must be identical at each period t. Players in BASE2 face the aggregated choices of players in each other hat, at each period t. To see how this aggregation might be helpful, imagine the case where $m = 1$, and player A observed that the other player under hat 1 had chosen a_b in the previous period. In the absences of the common knowledge condition, player A is unsure if the other player had done so because he had ascertained his hat to be B or had simply randomised. Now imagine the same situation involving player A with $m = 6$, where all players under hat 1 had chosen a_b in the previous period. Since players choose their actions independently and simultaneously, it is unlikely that all players pursing a randomisation strategy would end up with the same action. Thus player A should be more incline to reason that players under hat 1 must have ascertain their own hat colour to be B. Hence, increasing m might overcome some of problems associated with the lack of common knowledge conditions.

3.3 Equilibrium Analysis in TRADE

TRADE only differs from BASE2 on the availability of the pre-game stage. To show the equilibrium predictions in TRADE, we will first begin with the game stage and thereafter work backwards to the pre-game stage.

Players in TRADE enter the game stage with $x_{ij} \geq 1$ tokens. We know from BASE2, that the number of players in under each hat has no influence on the optimal choices. How about the token ownerships? The answer as it turns out is no. This is because if adherence to the optimal chocies is Pareto optimal for players with one token (as in BASE1 and BASE2), it must also be Pareto optimal for players with more than one token.

By backward deduction, players in the pre-game stage observing $b_i(s)$ should expect to ascertain their hats' colour in period $b_i(s)+1$ of the game stage. Given the token redemption structure, whatever colour it may be, players should hence expect the tokens to be redeemed at $\beta^*_{i_j} = \mu - b_i(s)\delta$ and by this logic, will only purchase additional tokens at prices $p_i \leq \mu - b_i(s)\delta$ or sell tokens at $p_i > \mu - b_i(s)\delta$. Since players only trade tokens with the other players under the same hat, this establishes the equilibrium price $p^*_i = \beta^*_{i_j} = \mu - b_i(s)\delta$, where players are indifferent between buying or selling tokens.

3.4 Equilibrium Payoffs

Given the equilibrium discussion in above sub-sections, the equilibrium payoff can be derived for players in each treatment by substituting p^*_i and $\beta^*_{i_j}$ where relevant

$$\Pi^*_{i_j} = \begin{cases} \beta^*_{i_j} = \mu - b_i(s)\delta & \text{if } 1_G = 0 \ \&\ x_{i_j} = 1 \\ p^*_i + (\beta^*_{i_j} - p^*_i)x_{i_j} = \mu - b_i(s)\delta & \text{if } 1_G = 1 \ \&\ x_{i_j} > 0 \\ p^*_i = \mu - b_i(s)\delta & \text{if } 1_G = 1 \ \&\ x_{i_j} = 0 \end{cases} \tag{2}$$

Notice that the equilibrium payoff ($\Pi^*_{i_j}$) only depends on $b_i(s)$ and is independent of the treatment variations. For any fixed n, the treatments are therefore payoff equivalent for any player observing $b_i(s)$. This is an important feature of this paper, since it allows us to motivate any potential differences between the respective treatments to the role of markets in the pre-game stage.

4 Behavioural Prediction

When players observe $b_i(s) = 0$, they should immediately ascertain their hats to be B. But when $b_i(s) > 0$, the process of ascertaining their hats' colour involve logical and epistemological reasoning. Thus players' sophistication are integral in the equilibrium discussions when $b_i(s) > 0$. If players are indeed heterogeneous in their sophistication, this paper presents a behavioural prediction that the pre-game stage in TRADE would "filter" out the unsophisticated players and only admit the sophisticated players into the game stage.

To see why this might be so, assume that the population of players under each hat consist on non-zero proportions of both Sophisticated types (players with sufficient abilities to know the optimal choices in the game stage) and unsophisticated types (players with insufficient abilities). When $b_i(s) > 0$, unsophisticated players will not expect to ever ascertain their true hat colours if they entered the game stage. As such, the optimal choice for such players upon entering the game stage would be to randomise between a_b and a_r in the very first period - choosing a_n is dominated for such player as he incurs a cost of δ with no expected revisions to his posterior in the later periods. His expected token redemption rate in the game stage is therefore $\mu - (1/2)\alpha$. As such, unsophisticated players should only purchase tokens at prices $p_i \leq \mu - (1/2)\alpha$ and sell tokens at prices $p_i > \mu - (1/2)\alpha$.

Assume for now that the sophisticated players always expects to ascertain their true hats' colour in the game stage. They should thus only purchase tokens at prices $p_i \leq \mu - b_i(s)\delta$ and sell his token at $p_i > \mu - b_i(s)\delta$. Since $(1/2)\alpha > (n+1)\delta$, at prices $p_i \in (\mu - (1/2)\alpha, \mu - b_i(s)\delta]$, it is therefore incentive compatible for sophisticated players to purchase tokens and unsophisticated players to sell tokens. Furthermore sophisticated players should know that given the availability of pre-game stage, the only players who will eventually enter the game stage must also be sophisticated players - this establishes common knowledge of sophistication. The behavioural prediction in this paper is therefore that at instances where players observed $b_i(s) > 0$, the markets in the pre-game stage of TRADE should result in the allocation of tokens to the sophisticated players.

5 Hypotheses Test

In the following discussion, we will omit the suffix i_j. In additional, we will make references to those instance where $b = 0, 1, 2$ black hats. The analysis will be performed in two steps, the first step compares the summary statistics between treatments and the second step employs econometric methodology to study the determinants of subjects' behaviours in the game stage.

To compare performances between treatments, one would naturally consider the *adherence rates* - the ratio of subjects in the game stage who were observed to have chosen actions consistent with the optimal choices. This is simply because adherence to the optimal choices is dependent on subjects' sophistication and results in their token being redeemed at the Pareto optimal rate $\beta^* = \mu - b\delta$. However, such a measure assigns uniform weights to the decisions of all subjects in the game stage. Since the idea in TRADE is for the allocation of tokens to the sophisticated players, comparison between treatments should provide some account for such allocation.

To account for token allocation, we first define a *efficiency event* to have occurred if a token belonging to a subject observing $b = 0, 1, 2$ was found to have been redeemed at $\beta^* = \mu - b\delta$. As such, three efficiency events are considered to have occurred if a subject with three tokens was found to have adhered to the optimal choices. The *efficiency rate* - the ratio of tokens redeemed at $\beta^* = \mu - b\delta$, is the natural extension of this term. Notice that the efficiency rate will assign greater weights to the decisions of subjects with more tokens. This will of course be irrelevant for the BASE1 and BASE2 treatment since the adherence rates and efficiency rates must always be equivalent - subjects enter the game stage with exactly one token. However, this will not necessarily be true for TRADE, since subjects may trade tokens in the pre-game stage. Hence, the efficiency rate would be a more appropriate measure of performances in direct comparisons between treatments. This leads us to the following two hypotheses test

H1: The aggregated efficiency rate in BASE1 is no different from that of BASE2.

H2: The aggregated efficiency rate in TRADE is higher than those reported in BASE1 and BASE2.

Aggregated efficiency rates here refers to the observed efficiency rate over all b instances. The first test serves as an empirical warm-up, where we examine the marginal influences of increasing m on the aggregated efficiency rate. Building on this finding, we can thus proceed to our main test where we examine the behavioural prediction presented in this paper, such that the ability to trade tokens in the pre-game stage of TRADE would result in the allocation of tokens to the sophisticated players. In this is indeed true, we should therefore expect the aggregated efficiency rate in TRADE to be significantly different and higher than the other two treatments.

Building on the findings in H1 and H2, we will proceed to the second step of our data analysis, where we study the determinants of subjects' behaviours. In this step, we are only concerned with those subjects who had entered the game stage, and thus the emphasis will be likelihood to which subjects would adhere to the optimal choices. This allows us to revisit H2 and also examine the behavioural prediction put forth in this paper, that sophisticated players in TRADE would purchase additional tokens and at prices $p \in (\mu - (1/2)\alpha, \mu - b\delta]$. This leads us to the following hypotheses that will jointly evaluated

H3: The likelihood of an agreement for subjects in TRADE is increasing with his inventory of tokens, at instances where $b = 1, 2$.

H4: The likelihood of an agreement for subjects in TRADE is strictly higher for subjects who had purchased tokens at prices $p \in (\mu - (1/2)\alpha, \mu - b\delta]$ relative to subjects who had purchased tokens at $p \notin (\mu - (1/2)\alpha, \mu - b\delta]$ or had not purchased tokens, at instances where $b = 1, 2$.

As discussed in the behavioural predictions, the effects of the token allocation should only be evidential at instances where subjects observed $b = 1, 2$.

6 Experimental Procedures

Two experimental sessions, were conducted for each treatment. Each session had involved 18 inexperienced subjects, recruited on a first come basis from the undergraduate cohort at the University of Exeter, through the ORSEE (Greiner, 2004) software. Table 2 reports on the subjects' demography for each treatment, by the schools they were enrolled into - Economic students study at the Business School. Although subjects had no formal training in game theory, those with stronger background in economics, engineering, mathematic or physics may potentially have some advantage with logical and epistemological reasoning due to their background training. This will be controlled for in the econometric analysis.

The experiments were conducted with the Z-Tree (Fischbacher, 2007) software and employed non-neutral framing of the problem. This was introduced to aid subjects unfamiliar with abstract reasoning problems. Each session had consisted of one practice round and ten paying rounds, where subjects payoffs were denoted in the fictitious currency, ECU. The following payoff parameters were employed: $\mu = 950$, $\delta = 50$, $\alpha = 700$ and $\bar{L} = 6000$. Subjects' overall payoffs were determined as the average over all ten rounds and converted into cash at the exchange rate of 67ECU/£1 in

Table 2: Demographics of Subjects by Schools Enrolled

School	BASE1	BASE2	TRADE
Business School	16	13	23
Engineering, Mathematics & Physical Science	3	1	5
Humanities	9	6	0
Life & Environmental Science	4	3	2
Social Sciences & International Studies	1	10	6
Others	3	3	0
Total	36	36	36

the BASE1 and BASE2 treatments, and 100ECU/£1 in the TRADE treatments.[11] The average duration of the BASE1 and BASE2 sessions were 95 minutes, whilst the TRADE sessions were 130 minutes. In addition to their experimental earnings, subjects received a show-up fee of £5 in the BASE1 and BASE2 sessions, and £8 in the TRADE sessions. Including the show-up fees, the average cash earnings were £16.64, £16.91 and £16.12 in the BASE1, BASE2 and TRADE treatments, respectively. Before collecting their cash payments, subjects were required to complete the Cognitive Reflective Test (Frederick, 2005) and declare any prior familiarity with the decisional task in the game stage.[12] For efficient comparisons between treatments, two sequences of states ($s \in S$) were randomly generated prior to the experimental proper. This was introduced to ensure that at each round of the respective treatments, there were the same number of subjects who observed $b = 0, 1, 2$ black hats.

Prior to experiment proper, we conducted a pilot test on the software and the instructions. The pilot test had raised some interesting problems with the experimental design, which prompted us to

[11]The difference in exchange rates was to control for any potential income effect due to a higher show-up fee being paid in the TRADE sessions.

[12]The Cognitive Reflective Test involves three question that triggers the wrong "instinctive" answer. (Q1) A bat and a ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost? (Q2) If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? (Q3) In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

make minor modifications to the design of BASE1 and BASE2. In the following, we will first detail the modification made and thereafter the experimental procedures in the respective treatment.

6.1 Minor Modifications to BASE1 and BASE2

Our pilot session was based on the BASE2 treatment design. One issue that was raised during this session was that subjects were sometimes observed to be adhering to the optimal choices despite the fact that they were following some randomisation process - through their feedbacks. This might be a common problem in experimental research since we only observe their behaviours and not the logical consequences of their behaviours.

To overcome the chances that adherence to the predicted behaviours were purely coincidental, we include an “outside option” for subjects to discretely end the game stage in a manner that does not affect the equilibrium analysis of the game. In addition to the actions a_b, a_r and a_n, subjects could also choose the outside option with the action “Toss a Coin, I will never know (a_c)”. If the subjects chooses a_c, he leaves the game stage with a fixed cost of 250 ECU. In doing so, he assigns the computer to choose the action a_b or a_r on his behalf - with equal probability.[13] However, only the experimenter would know that the subject had chosen a_c, whereas all other subjects who observe that he had chosen either a_b or a_r as determined by the computer. The action a_c will always be dominated in the equilibrium analysis and does not influence the optimal choices. The expected token redemption rate with adhering to the predicted behaviour is $950 - 50b$, with choosing a_c at any period t is $700 - 50(t - 1)$, and randomising with either a_b or a_r for uncertain players is $600 - 50(t - 1)$. Thus for subjects who do not expect to ever ascertain his hat colour, the expected token redemption rate with choosing a_c in now strictly higher than choosing all other actions. The outside option was omitted from TRADE, since an equivalent outside option already exist, the ability to sell your token and avoid the game stage altogether.

[13]For example if subject A choose a_n in the first period and a_c in the second period - the computer had chosen a_b on his behalf, his token will be redeemed at the rate of 950-50-250=650 ECU. All other subjects would have observed that A had chosen a_b in the second period.

6.2 BASE1

Upon entering the experiment, subjects were allowed 40 minutes to read through the instructions (see Appendix A) and complete a questionnaire, testing their understanding of the experimental design. Thereafter, subjects were randomly paired with two other players into a *group* and remained within the same group for the duration of the experiment - total of 12 group.[14] At the start each round, subjects were randomly assigned to one of three hats and were presented with the other hat colours within their group. Subjects were also informed that *there is at least one black hat* and proceed directly into the game stage. To avoid confusion, the notion of tokens were omitted from the subjects instructions. The game stage proceed as we had discussed and each period had lasted a maximum of 240 seconds. At each period $t > 1$, subjects were presented on their computer screens the period $t - 1$ actions of all other subjects within their group. A limitation of the software design was such that subjects had to proceed through the periods together. This meant that subjects who had chosen the actions a_b, a_r or a_c were facing a blank screen as they waited for other subjects to proceed through the periods. However, subjects were observed to have taken the opportunity to "sketch" their behaviours in the game.

6.2.1 BASE2

The sessions differ from the BASE1 sessions in the following: Each group consisted of 18 subjects, with 6 subjects under each hat (see Appendix B for the instructions) However, subjects were again randomly assign to one of three hats in each round. At each period $t > 1$, subjects were presented on their computer screens a table that depicted the aggregated period $t - 1$ actions, by all subjects in the respective hats. For example, subjects under hat 1, will observe the relative frequencies of the actions a_b, a_r and a_n, chosen by all subjects under hat 2 and 3.

6.2.2 TRADE

Each group again consisted of 18 subjects with 6 subjects under each hat (see Appendix C for the instructions) When the round begins, subjects first observed the other hats' colours. Thereafter,

[14] This restriction was introduced since physical limitations allows us to only fit 18 subjects in each session. Thus subjects in BASE2 and TRADE will only interact with the same other 17 players for the duration of the experiment.

subjects enter the pre-game stage, where trade was facilitated through a continuous double auction (CDA) mechanism that lasted for 120 seconds - the market only consisted of the other subjects under the same hat.[15] Here, a price ceiling of 1200 ECU was imposed on the bid and ask prices, to restrict subjects from intentionally making losses. This also ensures that each subject was not capital constrained from purchasing all other tokens within his market.

After the pre-game stage had ended, only subjects with at least one token will enter the game stage - subjects without any tokens were able observe the proceedings of the game stage on their computer screens but prevented from participating. The game stage proceeded as described in the BASE1 sessions, with the exception that the action a_c was not available and the information available to subjects at each period $t > 1$. Here, their computerised screens depicted the aggregated period $t-1$ actions, by all subjects under the respective hats ranked by their token ownership. For example, subjects under hat 1, will observe the relative frequencies of the actions a_b, a_r and a_n chosen by those subjects under hat 2 and 3 with one, two, three,..., six tokens.

Since the loan of 6000 ECU had to be repaid at the end of the round, this implies that some subjects may have incurred negative payoffs - 20 observed bankruptcy out of the 360 instances. We thus introduced a lower bound of 0 ECU to restrict subjects from making negative payoffs in any round.

7 Results

Before we describe our experimental result, it is worth to remember that our experimental procedures ensure that there will be the same number of subjects starting each round in the respective treatments, observing $b = 0, 1, 2$ black hats. However, the ability to trade tokens meant that only a subset of subjects in TRADE would eventually enter the game stage. Nevertheless, there will still be the same number of tokens due for redemption at the rates $\beta^* = 950 - 50b$. In the following sub-sections, we will first present the summary statistics in all treatments to examine H1 and H2. Thereafter, we will focus on the prices in TRADE and give an overview to H3 and H4. Finally, we revisit H2 in our econometric analysis, where we will also consider H3 and H4.

[15] See Sunder (1995) for a survey on how the CDA markets might be efficient mechanisms in asset market experiments.

7.1 Summary Statistics

We present on Table 3, the efficiency rates in BASE1, BASE2 and TRADE. Each cell depicts the total number of tokens redeemed at the equilibrium rate $\beta^* = 950 - 50b$ (efficiency events) with the ratio in parenthesis. The final column of each panel depicts the aggregated efficiency rate for that round and the final row, over all rounds.

Interpretation of BASE1's and BASE2's results should be straightforward. For example in round 1 of BASE1, there were 24 subjects who began the round observing $b = 1$. Obviously all subjects entered the game stage. However, only 16 of those subjects were found to have adhered to the optimal choices, and thus only 16 tokens were redeemed at $\beta^* = 900 - 50(1) = 900$ ECU. The efficiency rate was thus computed as $16/24 \approx 0.67$. Interpretation of TRADEs results are less straightforward. In the round 1 of TRADE, there were again 24 subjects who began the round observing $b = 1$. However, after trading tokens in the pre-game stage, only 19 subjects had eventually entered the game stage. Out of these 19 subjects, 11 subjects were observed to have adhered to the optimal choices but 12 tokens were redeemed at the equilibrium rate - this implies that one of the 11 subjects must be owning two tokens. The efficiency rate was computed to be $12/24 = 0.5$. Pairwise comparisons between treatments will be made for each b instances with the Chi-Square test and Fisher's one-tail Exact test, where the p-values are reported as ρ and $\hat{\rho}$ respectively.

Let us first consider the results in BASE1 and BASE2. At instances where subjects observed $b = 0$, the efficiency rates in both treatments were unity. This should not be surprising, since given the public announcement that "*there is at least one black hat*", each subject should have immediately ascertained their hats to be black and choose a_b in the first period. This perhaps suggest that no subject had misunderstood the payoff structure. At instances where subjects observed $b = 1$, the optimal choices are less trivial or obvious and required the subjects to employ some logical reasoning. However, at most instances, the majority of subjects had understood the optimal choices. Here, the efficiency rates were observed to be 0.70 and 0.68 in BASE1 and BASE2 respectively and not significantly different ($\rho = 0.741$, $\hat{\rho} = 0.413$). At the most complex instances of the game stage, where subjects observed $b = 2$, the efficiency rates were now observed to be 0.14 and 0.15 in BASE1 and BASE2 respectively - the difference was again not found to be significant

Table 3: Efficiency Rates (BASE1, BASE2 and TRADE)

	BASE1				BASE2				TRADE			
Round	$b = 0$	$b = 1$	$b = 2$	Agg.	$b = 0$	$b = 1$	$b = 2$	Agg.	$b = 0$	$b = 1$	$b = 2$	Agg.
I	12(1.0)	16(.67)	-	**28(.78)**	12(1.0)	16(.67)	-	**28(.78)**	12(1.0)	12(.50)	-	**24(.67)**
II	12(1.0)	21(.88)	-	**33(.92)**	12(1.0)	18(.75)	-	**30(.83)**	12(1.0)	12(.50)	-	**24(.67)**
III	-	17(.71)	1(.08)	**18(.50)**	-	17(.71)	1(.08)	**18(.50)**	-	8(.33)	5(.42)	**13(.36)**
IV	-	7(.58)	4(.17)	**11(.31)**	-	8(.67)	3(.13)	**11(.31)**	-	4(.33)	2(.08)	**6(.17)**
V	-	7(.58)	2(.08)	**9(.25)**	-	7(.58)	5(.21)	**12(.33)**	-	8(.67)	5(.21)	**13(.36)**
VI	-	7(.58)	5(.21)	**12(.33)**	-	6(.50)	3(.13)	**9(.25)**	-	6(.50)	8(.33)	**14(.39)**
VII	6(1.0)	15(.63)	3(.50)	**24(.67)**	6(1.0)	15(.63)	1(.17)	**22(.61)**	6(1.0)	11(.46)	0(.00)	**17(.47)**
VIII	-	18(.75)	1(.08)	**19(.53)**	-	18(.75)	2(.17)	**20(.56)**	-	14(.58)	1(.08)	**15(.42)**
IX	12(1.0)	20(.83)	-	**32(.89)**	12(1.0)	18(.75)	-	**30(.83)**	12(1.0)	10(.42)	-	**22(.61)**
X	-	6(.50)	2(.08)	**8(.22)**	-	8(.67)	4(.17)	**12(.33)**	-	10(.83)	2(0.08)	**12(.33)**
Agg.	42(1.0)	134(.70)	18(.14)	**194(.54)**	42(1.0)	131(.68)	19(.15)	**192(.53)**	42(1.0)	95(.49)	23(.18)	**160(.44)**

($\rho = 0.859$, $\hat{\rho} = 0.500$). The fall in efficiency rates between instances at $b = 1$ to $b = 2$ is fairly obvious. This suggest that the decisional task at instances where $b = 2$ might be too complicated for most subjects. This is evidential in their decisions, where 50% and 40% of the observations in BASE1 and BASE2 respectively, resulted in subjects deviating at the very first period of the game stage. Finally, the aggregated efficiency rates over all rounds and b instances were observed to be 0.54 and 0.53 in BASE1 and BASE2 respectively. This was again not found to be significantly different ($\rho = 0.881$, $\hat{\rho} = 0.470$)

Result 1: *Consistent with H1, the aggregated efficiency rates in BASE1 and BASE2 were not found to be significantly different.*

This some extend, result 1 is convenient since it suggest that increasing the number of subjects under each hat has little or no obvious influences on their decisions in the game stage. Therefore, if any obvious differences were observed in TRADE, we could likely attribute them to the introduction of the pre-game stage.

We now turn to the central focus of this paper, the summary statistics in TRADE. At instances where subjects observed $b = 0$, the efficiency rate in TRADE was unity. However, at instances where subjects observed $b = 1$, the efficiency rate in TRADE was now observed to be 0.49, significantly lower and different to those reported in BASE1 and BASE2 ($\rho < 0.001$ and $\hat{\rho} < 0.001$ in all comparisons). At instances where subjects observed $b = 2$, the agreement frequency in TRADE was observed to be 0.18. This might seem higher than those reported in BASE1 and BASE2, but the difference was not found to be significant ($\rho > 0.393$ and $\hat{\rho} > 0.495$ in all comparisons). Finally, the aggregated efficiency rate in TRADE was found to be 0.44, significant lower than those reported in BASE1 and BASE2 ($\rho < 0.018$ and $\hat{\rho} < 0.012$ in all comparisons).

Result 2: *Contrary to H2, the aggregated efficiency rate in TRADE was found to be significantly lower than those in BASE1 and BASE2.*

This result suggest that allowing subjects to trade tokens in the pre-game stage had actually exacerbated the deviations from the optimal choices in the game stage, relative to the control treatments. Furthermore, comparisons between treatments suggest that such differences were primarily attributed to instances in TRADE where subjects observed $b = 1$ black hats. How might we reconcile such discrepancies? We suspect that this to be symptomatic of the complexity in the decisional task. When $b = 0$, the task was too trivial, and thus we do not observe any differences between the treatments. When $b = 2$, the task was too complex for most subjects, thus any marginal influence from the ability to trade tokens were minimal. As such, the "tipping point" lies at instances where subjects observed $b = 1$. This raises the question as to why the efficiency rates in TRADE might be lower than those in BASE1 and BASE2. This will explored in greater details in the following subsections.

With repeated games, the reader might be concerned with potential learning over rounds. We find little evidence of learning. The aggregated efficiency rates over rounds I-V were found to be 0.55, 0.55 and 0.44 in BASE1, BASE2 and TRADE respectively. The same rates over rounds VI-X were found to be 0.53, 0.52 and 0.44 respectively.

7.2 Prices and Decisions of Subjects in TRADE

A plausible explanation to Result 2, could be that the pre-game stage had often resulted in the tokens being purchased by the unsophisticated subjects. If this explanation is to be supported, we should first observe tokens in the pre-game stage of TRADE to be frequently transacted over the prices $p > 950 - 50b$. The rationale here is that at such prices, even sophisticated players would strictly benefit by selling their tokens and avoiding the game stage altogether.[16]

We present on Figure 1, the mean transaction prices in markets of the pre-game stage in TRADE, where subjects had observed $b = 0, 1, 2$ black hats. The horizontal line indicates the equilibrium price $p^* = 950 - 50b$ in those markets. The weighted volume of trade was found to be 1.11, 0.80 and 0.82 (total number of trades as a ratio of the the total number of tokens available for trade) in markets where subjects observed $b = 0, 1, 2$ respectively. Given that each market only consist of six subjects, there still seems to be robust numbers of transactions in the respective markets.

[16] Recall from the equilibrium analysis that adherence to the optimal choices results in each token being redeemed at the Pareto optimal rate $\beta^* = 950 - 50b$. Since, purchase of tokens are finance by loans, sophisticated players should only benefit from purchasing additional tokens at prices $p < 950 - 50b$.

Figure 1: Mean Prices in Pre-Game Stage (TRADE)

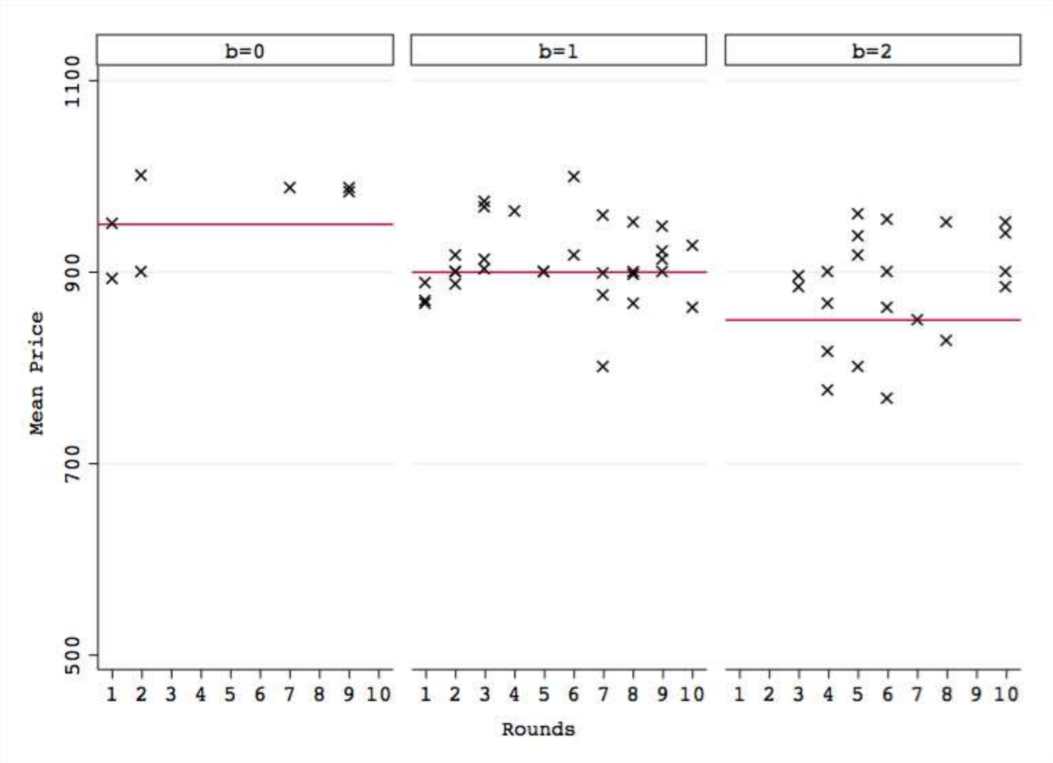

Mean prices were frequently found to be above the predicted trade prices - 57%, 46% and 71% of observations in markets where subjects observed $b = 0, 1, 2$ respectively - which are indicative of mis-pricing activities in the markets. Such mis-pricing or "bubbles" are a common occurrence in experimental markets (see Noussair and Tucker, 2013; Palan, 2013, for recent surveys on experimental markets). Clearly at instances where $b = 0$, price bubbles had no implications on the decisions of subjects in the game stage. Could the price bubbles at instances where subjects observed $b = 1, 2$ have led to unsophisticated subjects purchasing tokens? If this is the case, we should therefore expect to observe that subjects in TRADE with more tokens are proportionately less frequently found to have adhered to the optimal choices, relative to subjects with less tokens.

We present on Table 4 the adherence rates in TRADE by token ownership. For example, there were 70 instances where subjects observing $b = 1$ had entered the game stage with exactly one token, out of which 47 resulted in the subjects adhering to the optimal choices - the adherence rate was therefore $47/70 \approx 0.67$. At instances where $b = 0$, the ownership of tokens has no influence

Table 4: Adherence Rates by Token Ownership (TRADE)

Tokens	$b = 0$	$b = 1$	$b = 2$	Agg.
1	10(1.0)	47(.67)	12(.32)	69(.59)
2	9(1.0)	16(.47)	1(.05)	26(.41)
3	3(1.0)	4(.33)	3(.33)	10(.42)
4	-	1(.50)	0(.00)	1(.17)
5	1(1.0)	0(.00)	-	1(.33)
6	-	-	0(.00)	0(.00)
Agg.	23(1.0)	68(.57)	16(0.23)	107(.50)

on the adherence rates. However, at instances where $b = 1$ or $b = 2$, the adherence rates seem to decrease with token ownership. This is consistent with the explanation that price bubbles in the markets had often led to the unsophisticated subjects purchasing additional tokens.

A possible criticism of our analysis is such that we have implicitly assumed that subjects' behaviours in the game stage are correlated to the prices which they had purchased additional tokens. The behavioural prediction of this paper suggest that sophisticated players should be purchasing tokens at prices $p \in (600, 950 - 50(b)]$ at instances where $b > 0$. To provide some insights on this matter, we derived for each subject in, his average purchase price ($\ddot{p}$), which was computed as the sum of all his purchasing expenditure in the pre-game stage divided by the total number of tokens purchased.[17]

We present on Figure 2 the average purchase price (horizontal axis) and adherence to the optimal choices (vertical axis). Each observation indicates the average purchase price for an individual subject and whether he was found to have adhered to the optimal choices in the game stage (the numeral 1 indicates that the subjects had adhered). This of course excludes all observation where subjects were *inactive* - did not purchase tokens in the pre-game stage - or had sold all their tokens.

It is difficult to see any linear relationship between adherence and $\ddot{p}$ and there is no theoretical

[17] As trade was facilitated through a continuous double auction mechanism, subjects could purchase and sell token simultaneously within the trading period. Thus the average purchase price seeks to normalise his overall purchasing activities within the trading period. One could alternatively consider the average sale price, however we prefer to work with the purchasing activities since it may better describes a subjects expected token redemption rate.

Figure 2: Average Purchase Price ($\ddot{p}$) and Adherence to Optimal Choices (TRADE)

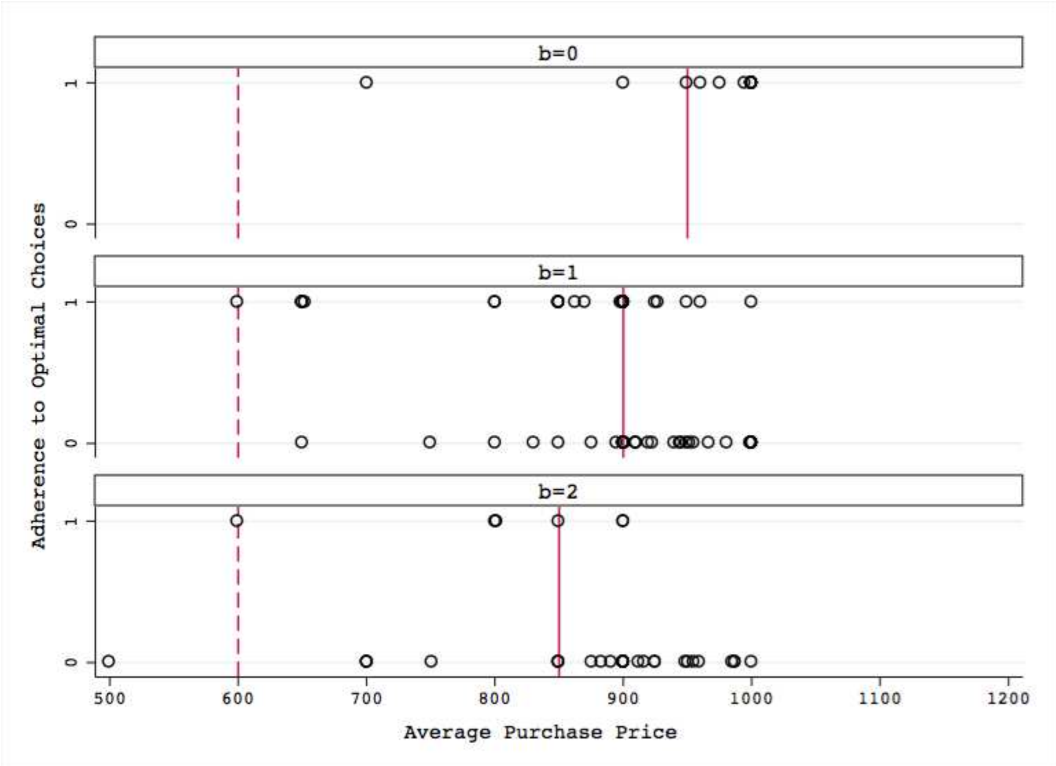

Table 5: Adherence Rates by Average Purchase Price (TRADE)

	$b = 1$	$b = 2$
$\ddot{p} \in (600, 950 - 50(b)]$	27(0.67)	3(0.30)
$\ddot{p} \notin (600, 950 - 50(b)]$	6(0.22)	3(0.11)

justification for one. We can partition the observations into two (crude) clusters, those with $\ddot{p} \in (600, 950-50(b)]$ and $\ddot{p} \notin (600, 950-50(b)]$. In doing so, we notice some relationship between $\ddot{p}$ and decisions of subjects in the game stage.To see this more clearly, we present on Table 5 the adherence rates at instances where $\ddot{p} \in (600, 950-50(b)]$ and $\ddot{p} \notin (600, 950-50(b)]$. One immediately observes the rates to be higher in the former relative to the latter condition when $b = 1$ - 0.67 and 0.22 respectively. However, when $b = 2$, the differences - 0.30 and 0.11 in the former and latter cases respectively - do not seem obvious.

Token together, the discussions so far seem consistent with the explanation that price bubbles in the pre-game stage of TRADE had led to mostly unsophisticated players buying tokens. We will revisit this explanation in the econometric analysis.

7.3 Econometric Analysis

This section employs econometric methods to re-examine the previous sections' results, and jointly investigate H3 and H4, whilst controlling for subject specific characteristics. Given that subjects remained within the same group for the duration of the experiment, one should hence expect the residual estimates to be highly correlated amongst subjects of the same group but independent from those of other groups. As such, the approach taken in this paper follows that of Bayer and Chan (2007), with the three-level hierarchical Generalised Linear Latent and Mixed Model (Rabe-Hesketh et al. (2005)).

The first level refers to observations at round r, the second level refers to subjects indexed by l and the third level refers to groups indexed by g. To ensure variations in the data, all observations where subjects observed $b = 0$ or when subjects did not participate in the game stage were excluded. This results in 827 level-1 variables, 105 level-2 variables and 16 level-3 variables.[18] The regression

[18] Although there were a total of 108 subjects in all treatments, there were three subjects in the TRADE treatment

model adopts a *logistic* link function

$$\text{Logit}[\text{Prob}(y_{rlg} = 1)|\mathbf{x}_{rlg}, \zeta_{lg}^{(2)}, \zeta_{g}^{(3)}] = \mathbf{x}'_{rlg}\beta + \zeta_{lg}^{(2)} + \zeta_{g}^{(3)} \tag{3}$$

where the dependent variable refers denotes the adherence to the optimal choice in round r, by subject l belonging to crowd g. The model assumes that $\zeta_{lg}^{(2)}|\mathbf{x}_{rld}, \zeta_{g}^{(3)} \sim \mathcal{N}(0, \psi^{(2)})$, where $\psi^{(2)}$ denotes the between-subject, within-group variances. Furthermore, it assumes that $\zeta_{g}^{(3)}|\mathbf{x}_{rlg} \sim \mathcal{N}(0, \psi^{(3)})$, where $\psi^{(3)}$ denotes the between-group variance. For the purposes of this paper, the observations from BASE1 and BASE2 were pooled together to form the BASE observations. Thereafter interactive dummies were introduced to distinguish TRADE observations from those of BASE at instances where subjects observed $b = 1$ and $b = 2$ - these are dummy variables in our regressions. Five regression models were considered, where each estimation process employs the adaptive quadrature numerical methods to maximise the marginal log-likelihoods.[19] The estimations results are reported on Table 6, where the p-values are presented in parenthesis. The likelihood-ratio test prefers regression 4 to all other regressions (at the 1% significance level) but our discussions will make references to regression 5, since it introduces some subject specific coefficients. The estimates in the respective regression were also found to be consistent to the random-effects logistic regression results - not reported here.

The discussion henceforth will make references to the average subject, a hypothetical subject where the coefficient estimates are set to their averages. We will revisit H2 to investigate if an average subject in TRADE is indeed significantly less likely to adhere to the optimal choices, relative to an average subject in BASE. The log-likelihood of adherence decreases by 1.38 at instances where $b = 1$, and increases by 0.69 at instances where $b = 2$, for an average subject in TRADE relative to an average subject in BASE. However, only the former was found to be mildly significant (p-value=0.085).

Result 2': *Consistent with Result 2, the likelihood of an adherence for subjects in TRADE was significantly lower at instances where $b = 1$. At instances where $b = 2$, no significant effect was*

who always sold their tokens when they observed $b = 1$ or $b = 2$. There were hence only 105 level-2 variables in the regression.

[19]The adaptive quadrature method was employ to increase computation efficiency and estimation precision (Rabe-Hesketh et al. (2002)).

Table 6: Econometric Regression Results: Generalised Linear Latent and Mixed Model

Dependent Variable: Adherence to the Optimal Choices					
Regression	(1)	(2)	(3)	(4)	(5)
Coefficient	Est.	Est.	Est.	Est.	Est.
$(b=2)$	-3.911^{***} (0.000)	-4.327^{***} (0.000)	-3.936^{***} (0.000)	-4.810^{***} (0.000)	-4.831^{***} (0.000)
$(b=1)\times$TRADE	-0.936^{*} (0.065)	-0.567 (0.291)	-1.413^{**} (0.010)	-1.010^{*} (0.081)	-1.381^{*} (0.085)
$(b=2)\times$TRADE	0.999^{*} (0.096)	1.089^{*} (0.093)	0.943 (0.140)	1.080 (0.106)	0.695 (0.313)
$(b=1)\times$Token	-	-0.610^{*} (0.062)	-	-1.111^{***} (0.004)	-1.131^{***} (0.007)
$(b=2)\times$Token	-	-0.169 (0.723)	-	-0.206 (0.676)	-0.201 (0.681)
$(b=1)\times$TRADE$\times\ddot{p}\in(600,950-50(b)]$	-	-	1.399^{**} (0.012)	2.016^{***} (0.002)	2.015^{***} (0.002)
$(b=2)\times$TRADE$\times\ddot{p}\in(600,950-50(b)]$	-	-	0.162 (0.860)	0.018 (0.984)	0.029 (0.976)
Sequence	-	-	-	-	0.252 (0.570)
Familiarity	-	-	-	-	-0.622 (0.314)
Gender	-	-	-	-	0.676 (0.147)
CRT Score	-	-	-	-	0.054^{*} (0.09)
Business School	-	-	-	-	1.029 (0.285)
Eng, Math & Phy Science	-	-	-	-	1.852 (0.128)
Humanities	-	-	-	-	0.841 (0.431)
Life & Environmental Science	-	-	-	-	1.111 (0.343)
Social Science & Int'l Studies	-	-	-	-	0.922 (0.375)
Constant	1.194^{***} (0.000)	1.798^{***} (0.000)	1.201^{***} (0.000)	2.303^{***} (0.000)	1.054 (0.293)
$\psi^{(2)}$	3.68^{***}	3.54^{***}	3.82^{***}	3.64^{***}	3.26^{***}
$\psi^{(3)}$	$\frac{1.7}{10^{11}}$	$\frac{2.8}{10^{10}}$	$\frac{2.4}{10^{11}}$	$\frac{4.8}{10^{11}}$	$\frac{1.2}{10^{12}}$
Negative Log-Likelihood	395.01	393.12	391.62	387.51	384.20

***:p-value<0.01; **:p-value<0.05 and *:p-value<0.10

Level-1 Observations $n=827$; Level-2 Observations $n=105$ and Level-3 Observations $n=16$.

observed.

For the average subject in TRADE, the log-likelihood of adherence decrease by 1.13 when $b = 1$ and 0.20 when $b = 2$, for each additional token owned. Again, only the former coefficient was again found to be significant (p-value=0.007). This is consistent with the findings on Table 4, where adherence rates in TRADE were observed to have decrease with token ownership. However, the surprise here is such that the coefficients were only significant at instances where subjects observed $b = 1$.

Finally, for the average subject in TRADE, the log-likelihood of agreement increases by 2.015 and 0.029 at instances where $b = 1$ and $b = 2$ respectively, when the subjects average purchase price was $\ddot{p} \in (600, 950 - 50(b)]$, relative to other instances when the subject was found to be inactive or $\ddot{p} \notin (600, 950 - 50(b)]$. Once again, only the former was found to be significant (p-value=0.002). To confirm these findings, we also considered an alternative regression where the interactive dummy variables were specified for subjects in TRADE who were observed to have purchased tokens at prices $\ddot{p} \in (600, 950 - 50(b)]$. Here the likelihood of adherence was found to be significantly lower for subjects with $\ddot{p} \in (600, 950 - 50(b)]$ at instances where $b = 1$ but not significantly different at instances where $b = 2$.

This econometric results suggest that after controlling for purchases prices and token ownership, the likelihood of adherence of an average subject in TRADE is mildly different to an average subject in BASE at instances where $b = 1$ and significantly different when $b = 2$. Thus the determinacy of differences in the aggregated efficiency rates reported in the previous sections were predominantly driven by subjects who were purchasing tokens at elevated prices. This brings us to the following results

Result 3: *Contrary to H3, the likelihood of adherence to the optimal choices was decreasing with token ownership for subjects in TRADE at instances where $b = 1$. At instances where $b = 2$, no significant effect was observed.*

Result 4: *Consistent with H4, the likelihood of adherence to the optimal choices was found to higher for those whose $\ddot{p} \in (600, 950 - 50(b)]$ relative to those who were inactive or whose $\ddot{p} \notin$*

$(600, 950 - 50(b)]$ *at instances where subjects in TRADE observed* $b = 1$*. At instance where* $b = 2$*, no such relationships were found to be significant.*

We did not find any significant effects due to differences in gender, sequence administered, schools or prior familiarity. The latter point is interesting since the decisional task for subjects in the game stage might be trivial if he had prior familiarity with the problem. However, this finding might highlight a central feature of the game stage, such that prior familiarity might only helpful if it was common knowledge. There is some mild evidence that the likelihood of adherence is increasing with the subjects' scores in the Cognitive Reflective Test (CRT). The CRT test involved three questions which required subjects to employ some effort in thought and reasoning before providing the answers. Given that the adherence to the predicted behaviours in the game stage also requires effort and logical reasoning, perhaps the CRT test score is capturing some of these abilities. Despite the subject characteristics controls, there still seem to be significant between-subject variances ($\psi^{(2)}$) although the between crowd variance ($\psi^{(3)}$) was not found to be significant.

Once again our results raises the question as to why any differences between the BASE and TRADE treatments or within the TRADE treatment, were only found to be significant at instances where $b = 1$. Our prior on this matter is that the decisional task where $b = 2$ was too complicated or complex for most subjects to comprehend. This is evidential in the above regressions, where the log-likelihood of adherence was found to decrease by 4.831 (p-value=0.0001) in all treatments when a subject observes $b = 2$ relative to observing $b = 1$. By this extension, subjects may have perceived the predicted behaviour at instances where $b = 1$ to be similar to that where $b = 2$. If such misperception were indeed reflected in the purchase prices of tokens in TRADE, we should expect the average purchase prices at instances where $b = 1$ to not be significantly different from those where $b = 2$. We thus conducted a linear regression on the average purchase price ($\ddot{p}$) with the situation dummies b_0 and b_2 which refer to these instances where subjects observed $b = 0$ and $b = 2$, with $b = 1$ as the reference. The p-values are again reported in parenthesis.

$$\ddot{p} = \underset{(0.001)}{80.76} b_0 - \underset{(0.276)}{21.38} b_2; \quad N = 121,\ R^2 = 0.10 \tag{4}$$

The regression result found $\ddot{p}$ to be significantly higher at instances where $b = 0$ relative to instances where $b = 1$. However, at instances where $b = 2$, $\ddot{p}$ was not found to be significantly different from those at $b = 1$. This lends some weight to the possibility that subjects may have misperceived the optimal choices at instances where $b = 2$ to be similar to those where $b = 1$.

8 Conclusion

We began this paper with the interest of investigating the conventional wisdom that markets should result in the allocation of the rights for performing decisional tasks to those players who are best suited to perform the tasks. To do so we embedded the decisional tasks in a game motivated by Littlewood (1953), Red Hat puzzle. Three treatments were considered, BASE1, BASE2 and TRADE, where only in the last treatment were players permitted to trade their rights (in the form of tokens) to perform the decisional task of the game. The equilibrium analysis suggested that players' decisiosns in the game should be independent of the treatment variations, but we provided a behavioural prediction that the markets in the pre-game stage of TRADE would "filter" out unsophisticated players and allocate tokens to the sophisticated players.

In comparisons between treatments, we found the aggregated efficiency rate in TRADE - our measure of performance - to be significantly lower than those reported in BASE1 and BASE2. The difference was primarily driven by the decisions of subjects in TRADE at instances where they observed $b = 1$. For some insights to our results, we studied the transaction prices of tokens in the pre-game stage of TRADE. Here we found evidence of price bubbles, where mean prices were often found to be above the equilibrium price in markets where subjects observed $b = 0, 1, 2$ black hats. A such prices, even sophisticated players would strictly prefer to sell their tokens and avoid the game altogether. This suggest that the only players who would be purchasing tokens at such elevated prices must be the unsophisticated players. Further support for this explanation was provided in our econometric analysis, where we show that at instances where subjects observed $b = 1$, those who had purchased additional tokens were significantly less likely to be adhere to the optimal choices in the game. However, we also found that subjects who had on average purchase tokens at prices $\ddot{p} \in (600, 950 - 50(b)]$ at instances where $b = 1$, were significantly more likely to be in adherence. This is consistent with our prior that prices are an important determinant in the allocation of rights for decisional tasks. No significant effects were found for instances where $b = 0$ and $b = 2$.

The question therefore is why might subjects in TRADE be willing to purchase tokens at prices $p > 950 - 50b$, especially at instances where $b = 1$. Some possible explanation are offered here.

(i) The continuous double auction mechanism may induce speculative activities in the market, where subjects might be seeking to churn prices within the trading durations, often resulting with the unsophisticated subjects being stucked with the hot potato (Tirole, 1982).

(ii) Unsophisticated subjects may not be aware of their own limitations or overconfident (Odean, 1998; Shleifer, 2000) of their own abilities and hence mis-priced the tokens.

(iii)

Our results also shed some light on Kluger and Wyatts (2004) experiment findings with the Monty Hall problem, that when there were at least two sophisticated players in the market, prices will converge to the equilibrium price. Clearly this is not the case in our results, even when the summary statistics suggest there to be more than two sophisticated subjects. We conjecture that their results might have been due to the nature of the Monty Hall problem, where the instinctive price (unsophisticated price) is naturally lower than the equilibrium price. When the instinctive price is less obvious, such as in our paper, they results may no longer hold.

To some extend, our results might also shed some light to the empirical literature in corporate governance. Here, the act of purchasing another players' token can be viewed as a corporate takeover. In an extensive survey of the corporate takeover literature by Martynova and Renneboog (2008), the authors found little evidence that operating performances of the acquired firms had improved ex-post the takeover. Surveys on behavioural finance also suggest that bidding firms often overpay in corporate takeovers, a phenomenon usually known as the "Winner's Curse" (Thaler, 1988; Barberis and Thaler, 2003). Our experiments capture some of these observations as most subjects in TRADE who had purchased additional tokens, had done so at elevated prices and were found to have performed worser in the decisional task.

To conclude, our paper provides evidence that introducing a market where rights for performing tasks can be traded, do not naturally lead to the allocation of such rights to those players who might be best suited to perform the task. This is contradictory to the conventional wisdom and has important implications for any economic designer considering the best mechanism to allocate decisional task. Nevertheless, we see potential for such a design in other more straightforward games, e.g., Guessing Game (Nagel, 1995) , Centipede Game. This will be an ambition for further research.

References

Aumann, Robert J., 1976, Agreeing to disagree, *The Annals of Statistics* 4, 1236–1239.

Aumann, Robert J., 1999, Interactive epistemology I: Knowledge, *International Journal of Game Theory* 28, 263–300.

Barberis, Nicholas, and Richard Thaler, 2003, Chapter 18 a survey of behavioral finance, in M. Harris G.M. Constantinides, and R.M. Stulz, eds., *Financial Markets and Asset Pricing*, volume 1, Part B of *Handbook of the Economics of Finance*, 1053 – 1128 (Elsevier).

Bayer, Ralph C., and Mickey Chan, 2007, The dirty faces game revisited, School of Economics Working Papers 2007-01, University of Adelaide, School of Economics.

Brereton, J.M, 1986, *The British Soldier: A Social History from 1661 to the Present Day* (The Bodley Head Limited: London).

Bruce, Anthony P.C., 1980, *The Purchase System in the British Army, 1660-1871* (Royal Historical Society: London).

Camerer, Colin F., 2003, *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press).

Costa-Gomes, Miguel A., Vincent P. Crawford, and Bruno Broseta, 2001, Cognition and behavior in normal-form games: An experimental study, *Econometrica* 69, 1193–1235.

Fagin, Ronald, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi, 1995, *Reasoning About Knowledge* (Cambridge: MIT Press).

Fischbacher, Urs, 2007, z-tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171–178.

Frederick, Shane, 2005, Cognitive reflection and decision making, *Journal of Economic Perspectives* 19, 25–42.

Fudenberg, Drew, and Jean Tirole, 1993, *Game Theory* (MIT Press, Cambridge).

Geanakoplos, John D., 1994, Common knowledge, in Robert Aumann, and Sergiu Hart, eds., *Handbook of Game Theory with Economic Applications*, volume 2, chapter 40, 1437–1496 (Elsevier).

Geanakoplos, John D., and Heraklis M. Polemarchakis, 1982, We can't disagree forever, *Journal of Economic Theory* 28.

Greiner, Ben, 2004, The online recruitment system orsee 2.0 - a guide for the organization of experiments in economics, Working Paper Series in Economics 10, University of Cologne, Department of Economics.

Jensen, Michael C., and Richard S. Ruback, 1983, The market for corporate control, *Journal of Financial Economics* 11, 5–50.

Kluger, Brian D., and Steve B. Wyatts, 2004, Are judgement errors reflected in market prices and allocations? experimental evidence based on the monty hall problem, *The Journal of Finance* 59, 969–997.

Littlewood, John E., 1953, *A Mathematician's Miscellany* (Meuthen and Co Ltd).

Martynova, Marina, and Luc Renneboog, 2008, A century of corporate takeovers: What have we learned and where do we stand?, *Journal of Banking and Finance* 32, 2148–2177.

Maschler, Michael, Eilon Solan, and Shmuel Zamir, 2013, *Game Theory* (Cambridge University Press).

Milgrom, Paul, and Nancy Stokey, 1982, Information, trade and common knowledge, *Journal of Economic Theory* 26.

Myerson, Roger B., 1997, *Game Theory: Analysis of Conflict* (Harvard University Press).

Nagel, Rosemarie, 1995, Unraveling in guessing games: An experimental study, *American Economic Review* 85, 1313–1326.

Noussair, Charles N., and Steven Tucker, 2013, Experimental research on asset pricing, *Journal of Economic Surveys* 27, 554–569.

Odean, Terrance, 1998, Volume, volatility, price and profit when all traders are above average, *The Journal of Finance* 53, 1887–1934.

Palan, Stefan, 2013, A review of bubbles and crashes in experimental asset markets, *Journal of Economic Surveys* 27, 570–588.

Rabe-Hesketh, Sophia, Anders Skrondal, and Andrew Pickles, 2002, Reliable estimation of generalized linear mixed models using adaptive quadrature, *The Stata Journal* 2, 1–21.

Rabe-Hesketh, Sophia, Anders Skrondal, and Andrew Pickles, 2005, Maximum likelihood estimation of limited and discrete dependent variable models with nested random effects, *Journal of Econometrics* 128, 301 – 323.

Shleifer, Andrei, 2000, *Inefficient Markets* (Oxford University Press, Oxford).

Stahl, Dale O., and Paul W. Wilson, 1994, Experimental evidence on players' models of other experimental evidence on players' model of other players, *Journal of Economic Behavior and Organization* 25, 309–327.

Stahl, Dale O., and Paul W. Wilson, 1995, Models of other players: Theory and experimental evidence, *Games and Economic Behavior* 10, 218–254.

Sunder, Shyam, 1995, Experimental asset markets: A survey., in John H. Kagel, and Alvin E. Roth, eds., *Handbook of Experimental Economics*, 445–500 (Princeton University Press).

Thaler, Richard H., 1988, Anomalies: The winner's curse, *Journal of Economic Perspectives* 2, 191–202.

Tirole, Jean, 1982, On the possibility of speculation under rational expectations, *Econometrica* 50, 1163–1181.

Weber, Roberto A., 2001, Behavior and learning in the "Dirty Faces" game, *Experimental Economics* 4, 229–242.

A Instructions BASE1

There will be 10 experiment rounds, where you receive a payoff (denoted in ECU) at the end of each round. Upon completion of all 10 rounds, your earnings for the experiment will be computed as the average over ten rounds and converted to cash at the exchange rate of 67 ECU to £1. In addition, you will also receive a £5 show up fee. We shall now describe each experimental round.

Each round will involve three coloured hats - A, B and C. Each hat can be either Red or Black, and there is exactly one player under each hat.

1. You will not be able to see your own hat colour.

2. You will see on your computer screens the other two hat colours.

3. There will always be one black hat.

To determine the hat colours, the computer randomly picks from 1 of the 7 possible outcomes (see Table A1). or example, in outcome 2, Hat A is red, Hat B is black and Hat C is black. The player under Hat A will see that: Hat B is Black and Hat C is Black. The player under Hat B will see that: Hat A is Red and Hat C is Black. Finally, the player under Hat C will see that: Hat A is Red and Hat B is Black. There is an equal chance for any one of these outcomes. Notice that at each outcome, there will always be at least one black hat. See Figure A1 for an example of what you might observe. Here you are under Hat A and you see that both Hat B and C are black.

Your task in each round is to determine the colour of your hat. You will do this in the decision stage that will consist of 4 period. At each period the computer will present you with the following question, to which you must choose from 4 possible actions (a), (b), (c) or (d).

Computer's question: *"Do You Know your hat colour?"*

Your actions (a) My Hat is RED, (b) My Hat is BLACK, (c) No! I will decide in a later period and (d)Toss a Coin, I would never know.

Here are some rules:

RULE 1: *At each period, you have a maximum of 4 minutes to choose an action*

RULE 2: *You will immediately end the decision stage if the actions (a), (b) or (d) were chosen*

RULE 3: *You will only go to the next period if you had chosen (c) in the previous period*

RULE 4: *If you arrive at period 4, you can only chose from the actions (a), (b) or (d)*

RULE 5: *If you had chosen (d) the computer will simulate a coin toss and choose on your BEHALF either option (a) or (b) with equal chances*

RULE 6: *Any action chosen will be known to all other players in the subsequent period. Note: If you had chosen (d) and the computer chooses (b) on your behalf, the other players will only see that you had chosen (b).*

You are said to have "determined you hat colour" when you choose (a), (b) or (d). This is why you will only go to the next period if you have chosen (c). For example, if you had chosen (a) in period 1, the decision stage immediately ends for you. In period 2, all other players will observe that you had chosen (a) in period 1. However, if you had chosen (c) in period 1, you go on to period 2, when you must again choose your action. All players will also observe that you had chosen (c) in period 1. Here, are some screenshots to help you understand the decision stage design (Figure A2 and Figure A3).

Figure A2 presents an illustration of the first period in the decision stage. You are under Hat B and you see the other hat colours. In addition the computer presents you with the question "Do you know you hat colour" to which you must reply with one of the 4 possible actions.

Figure A3 presents an illustration of the second period in the decision stage. You are under Hat B and you see the other hat colours. The computer again presents you with the question "Do you know you hat colour" to which you must reply with one of the 4 possible actions. In addition, you also see on your screen the actions chosen by the other players in the previous period (period 1). You see that the player under Hat A had chosen (b) in period 1. You also see that the player under Hat C had chosen (b) in period 1. Finally, in this illustration you had chosen (c) in period 1.

After all players had ended the decision stage, your hat colour with be made known and your payoffs for the round will be determined. Your Payoffs depends on whether you had correctly determined you hat colour and the period which you had "determined your hat colour". If you had chosen (a) or (b), then your payoffs will depend on whether you are correct and the period which you had chosen them (see Table A2). If You had chosen (d), then your payoffs will only depend on the period which you had chosen (c) (see Table A3). Here are some examples to help you understand the payoffs:

The payoffs can be easily summarised as followed. You start the round with 950 ECU. You get 50 ECU deducted for each time you had chosen (c). In addition, you get 700 ECU deducted if you had chosen (a) or (b) and was found to be incorrect, or no deduction if found to be correct. If you had chosen (d), you'll get a fixed deduction of 250 ECU. Here are some examples to help you understand the payoff

(i) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (b). You have correctly determined your hat colour and your payoffs are therefore 850 ECU.

(ii) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (a). You have incorrectly determined your hat colour and your payoffs are therefore 150 ECU.

(iii) You hat is black. In period 1 you choose (c), in period 2 you (c) and in period 3 you choose (d). Your payoffs are therefore 600 ECU.

This completes the description of each experimental round. After the completion of 10 experiment rounds, we require you to complete a survey before you receive your cash payments. Please feel free to clarify any question or doubts you might have with regards to the instructions.

Table A1: BASE1: The 7 Possible Outcomes

Outcomes	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Chance	1/7	1/7	1/7	1/7	1/7	1/7	1/7
Hat A	B	R	B	B	B	R	R
Hat B	B	B	R	B	R	B	R
Hat C	B	B	B	R	R	R	B

B=BLACK and R=RED

Figure A1: Screen Shot (BASE1) - You see all other Hat colours

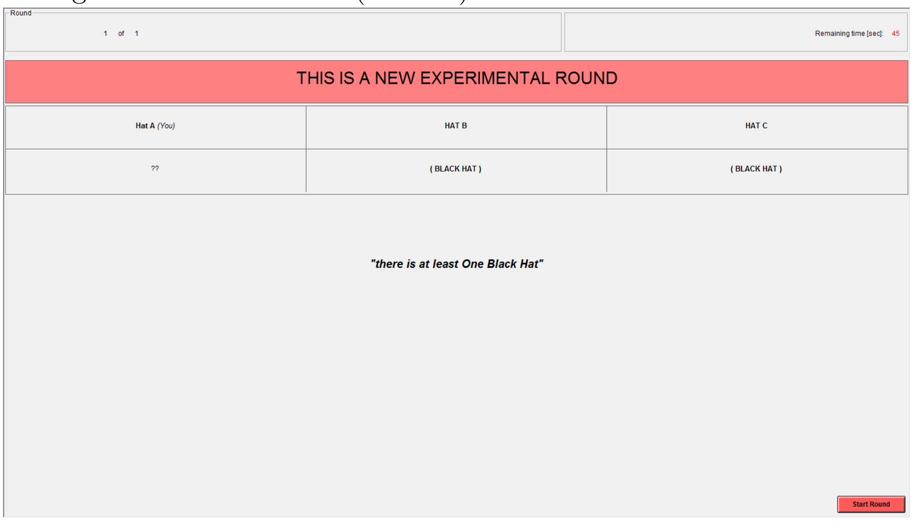

Figure A2: Screen Shot (BASE1) - Decision Stage Period 1

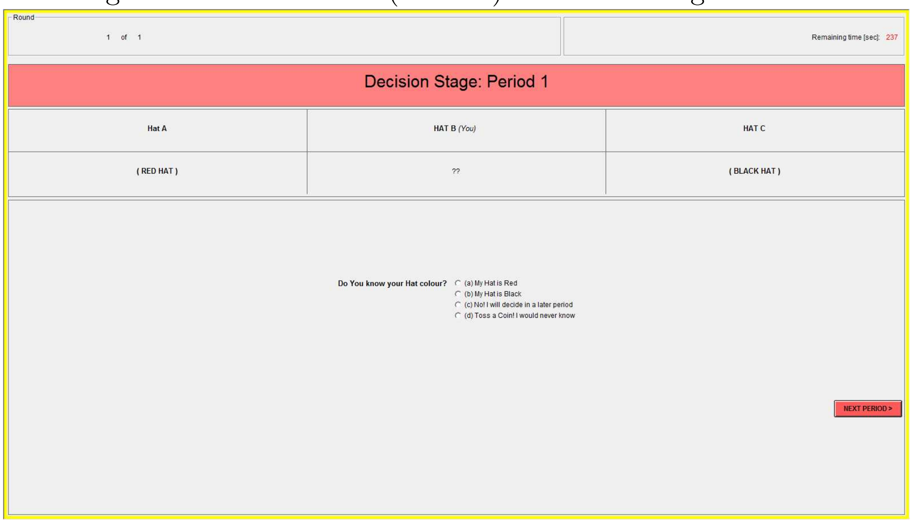

Figure A3: Screen Shot (BASE1) - Decision Stage Period 2

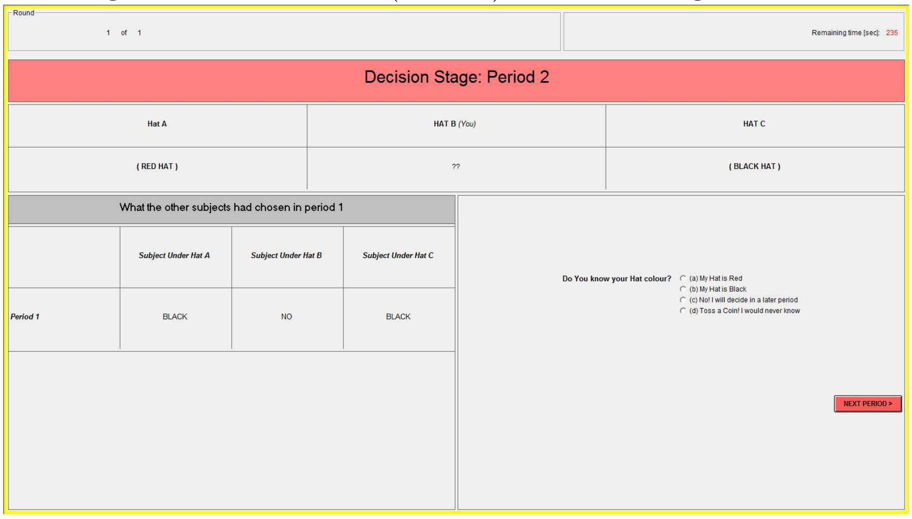

Table A2: BASE1: Payoffs with Choosing (a) or (b)

	Correct	Incorrect
"Determined Hat Colour" in Period 1	950 ECU	250 ECU
"Determined Hat Colour" in Period 2	900 ECU	200 ECU
"Determined Hat Colour" in Period 3	850 ECU	150 ECU
"Determined Hat Colour" in Period 4	800 ECU	100 ECU

Table A3: BASE1: Payoffs with Choosing (d)

"Determined Hat Colour" in Period 1	700ECU
"Determined Hat Colour" in Period 2	650ECU
"Determined Hat Colour" in Period 3	600ECU
"Determined Hat Colour" in Period 4	550ECU

B Instructions BASE2

There will be 10 experiment rounds, where you receive a payoff (denoted in ECU) at the end of each round. Upon completion of all 10 rounds, your earnings for the experiment will be computed as the average over ten rounds and converted to cash at the exchange rate of 67 ECU to £1. In addition, you will also receive a £5 show up fee. We shall now describe each experimental round.

Each round will involve three coloured hats - A, B and C. Each hat can be either Red or Black, and there is exactly six players under each hat.

1. You will not be able to see your own hat colour.

2. You will see on your computer screens the other two hat colours.

3. There will always be one black hat.

To determine the hat colours, the computer randomly picks from 1 of the 7 possible outcomes (see Table B1). or example, in outcome 2, Hat A is red, Hat B is black and Hat C is black. The players under Hat A will see that: Hat B is Black and Hat C is Black. The players under Hat B will see that: Hat A is Red and Hat C is Black. Finally, the players under Hat C will see that: Hat A is Red and Hat B is Black. There is an equal chance for any one of these outcomes. Notice that at each outcome, there will always be at least one black hat. See Figure B1 for an example of what you might observe. Here you are under Hat A and you see that both Hat B and C are black.

Your task in each round is to determine the colour of your hat. You will do this in the decision stage that will consist of 4 period. At each period the computer will present you with the following question, to which you must choose from 4 possible actions (a), (b), (c) or (d).

Computer's question: *"Do You Know your hat colour?"*

Your actions (a) My Hat is RED, (b) My Hat is BLACK, (c) No! I will decide in a later period and (d)Toss a Coin, I would never know.

Here are some rules:

RULE 1: *At each period, you have a maximum of 4 minutes to choose an action*

RULE 2: *You will immediately end the decision stage if the actions (a), (b) or (d) were chosen*

RULE 3: *You will only go to the next period if you had chosen (c) in the previous period*

RULE 4: *If you arrive at period 4, you can only chose from the actions (a), (b) or (d)*

RULE 5: *If you had chosen (d) the computer will simulate a coin toss and choose on your BEHALF either option (a) or (b) with equal chances*

RULE 6: *Any action chosen will be known to all other players in the subsequent period. Note: If you had chosen (d) and the computer chooses (b) on your behalf, the other players will only see that you had chosen (b).*

You are said to have "determined you hat colour" when you choose (a), (b) or (d). This is why you will only go to the next period if you have chosen (c). For example, if you had chosen (a) in period 1, the decision stage immediately ends for you. In period 2, all other players will observe that you had chosen (a) in period 1. However, if you had chosen (c) in period 1, you go on to period 2, when you must again choose your action. All players will also observe that you had chosen (c) in period 1. Here, are some screenshots to help you understand the decision stage design (Figure B2, Figure B3 and Figure B4).

Figure B2 presents an illustration of the first period in the decision stage. You are under Hat B and you see the other hat colours. In addition the computer presents you with the question "Do you know you hat colour" to which you must reply with one of the 4 possible actions.

Figure B3 resents an illustration of the second period in the decision stage. You are under Hat A and you see the other hat colours. The computer again presents you with the question "Do you know you hat colour" to which you must reply with one of the 4 possible actions. In addition, you also see on your screen the actions chosen by the other players in the previous period. For the six players under Hat A, all of them had chosen (c) in period 1. For the six players under Hat B, one of them had chosen (a), one of them had chosen (b) and four of them had chosen (c) in period 1. Finally for the six players under Hat C, two of them had chosen (a), one of them had chosen (b) and three of them had chosen (c) in period 1.

Figure B4 presents an illustration of the third period in the decision stage. You are under Hat A and you see the other hat colours. The computer again presents you with the question "Do you know you hat colour" to which you must reply with one of the 4 possible actions. For the six players under Hat A, all of them had chosen (c) in period 2. For the six players under Hat B, two of them had chosen (b) in period 2, two of them had chosen (c) in period 2 and two of them had not participated in period 2 since they had ended the round in period 1 and are awaiting results. For the six players under Hat C, three of them had chosen (c) in period 2 and three of them had not participated in period 2 as they had ended the round in an earlier period.

After all players had ended the decision stage, your hat colour with be made known and your payoffs for the round will be determined. Your Payoffs depends on whether you had correctly determined you hat colour and the period which you had "determined your hat colour". If you had chosen (a) or (b), then your payoffs will depend on whether you are correct and the period which you had chosen them (see Table B2). If You had chosen (d), then your payoffs will only depend on the period which you had chosen (c) (see Table B3). Here are some examples to help you understand the payoffs:

The payoffs can be easily summarised as followed. You start the round with 950 ECU. You get 50 ECU deducted for each time you had chosen (c). In addition, you get 700 ECU deducted if you had chosen (a) or (b) and was found to be incorrect, or no deduction if found to be correct. If you had chosen (d), you'll get a fixed deduction of 250 ECU. Here are some examples to help you understand the payoff

(i) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (b). You have correctly determined your hat colour and your payoffs are therefore 850 ECU.

(ii) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (a). You have incorrectly determined your hat colour and your payoffs are therefore 150 ECU.

(iii) You hat is black. In period 1 you choose (c), in period 2 you (c) and in period 3 you choose (d). Your payoffs are therefore 600 ECU.

This completes the description of each experimental round. After the completion of 10 experiment rounds, we require you to complete a survey before you receive your cash payments. Please feel free to clarify any question or doubts you might have with regards to the instructions.

Table B1: BASE2: The 7 Possible Outcomes

Outcomes	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Chance	1/7	1/7	1/7	1/7	1/7	1/7	1/7
Hat A	B	R	B	B	B	R	R
Hat B	B	B	R	B	R	B	R
Hat C	B	B	B	R	R	R	B

B=BLACK and R=RED

Figure B1: Screen Shot (BASE2) - You see all other Hat colours

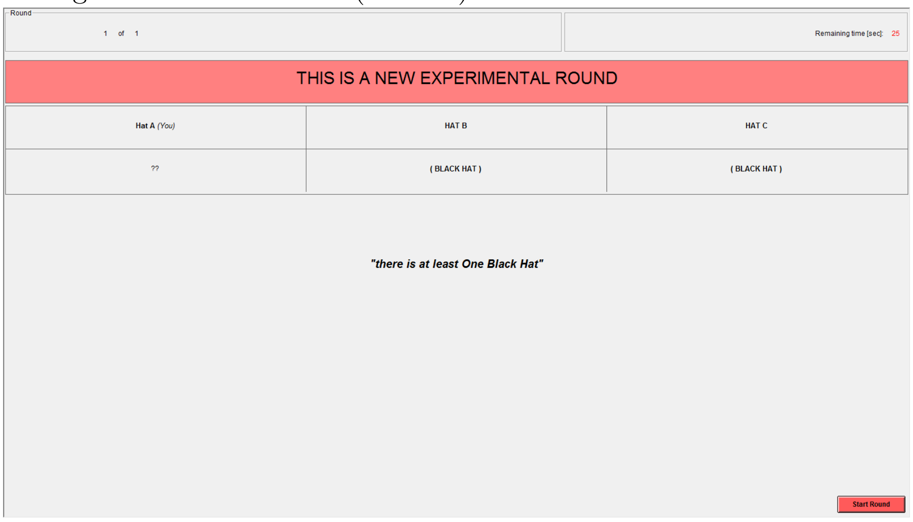

Figure B2: Screen Shot (BASE2) - Decision Stage Period 1

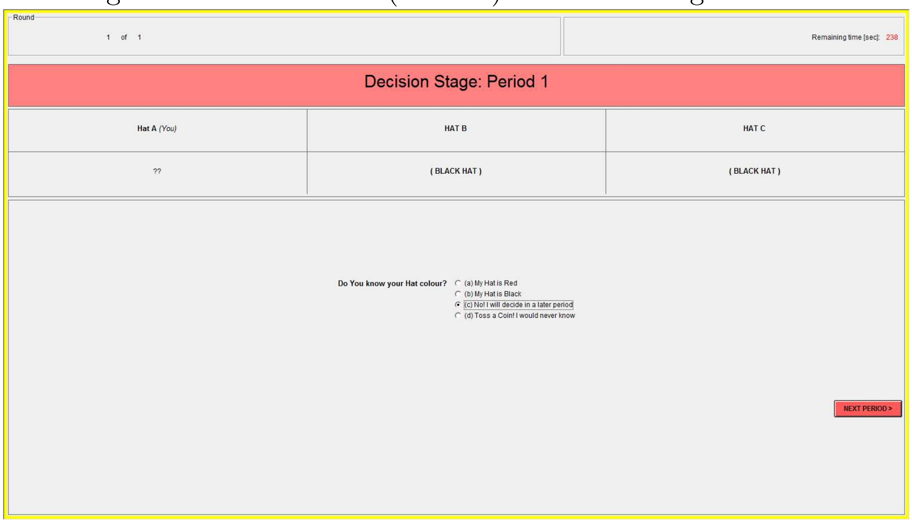

Figure B3: Screen Shot (BASE2) - Decision Stage Period 2

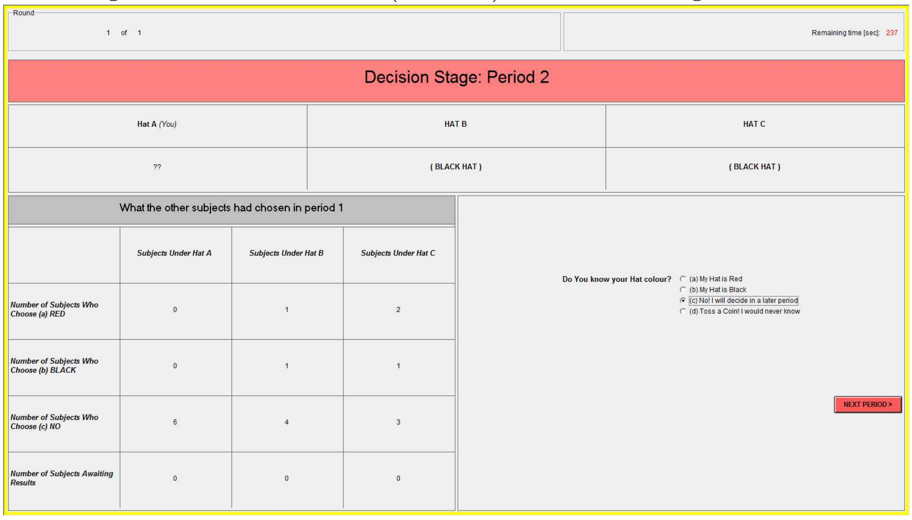

Figure B4: Screen Shot (BASE2) - Decision Stage Period 3

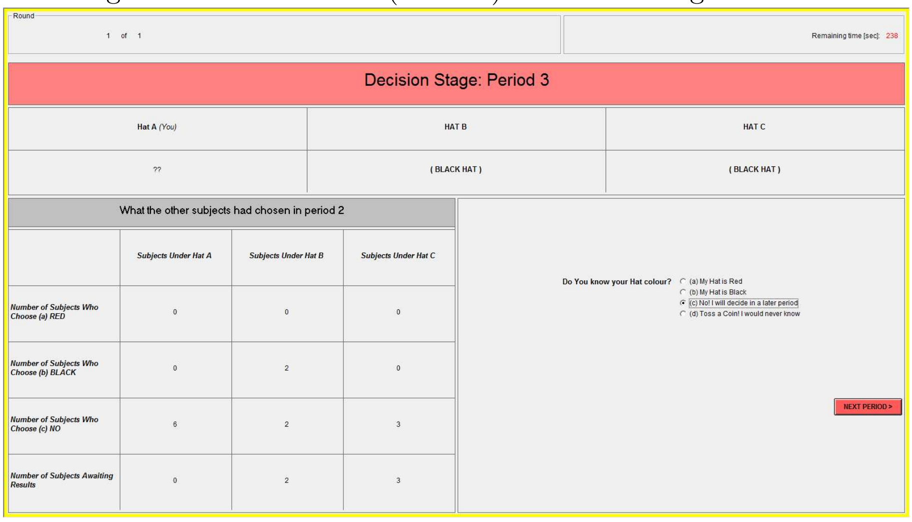

Table B2: BASE2: Payoffs with Choosing (a) or (b)

	Correct	Incorrect
"Determined Hat Colour" in Period 1	950 ECU	250 ECU
"Determined Hat Colour" in Period 2	900 ECU	200 ECU
"Determined Hat Colour" in Period 3	850 ECU	150 ECU
"Determined Hat Colour" in Period 4	800 ECU	100 ECU

Table B3: BASE2: Payoffs with Choosing (d)

"Determined Hat Colour" in Period 1	700ECU
"Determined Hat Colour" in Period 2	650ECU
"Determined Hat Colour" in Period 3	600ECU
"Determined Hat Colour" in Period 4	550ECU

C Instructions TRADE

There will be 10 experiment rounds, where you receive a payoff (denoted in ECU) at the end of each round. Upon completion of all 10 rounds, your earnings for the experiment will be computed as the average over ten rounds and converted to cash at the exchange rate of 100 ECU to £1. In addition, you will also receive a £8 show up fee. We shall now describe each experimental round.

Each round will involve three coloured hats - A, B and C. Each hat can be either Red or Black, and there is exactly six players under each hat.

1. You will not be able to see your own hat colour.

2. You will see on your computer screens the other two hat colours.

3. There will always be one black hat.

To determine the hat colours, the computer randomly picks from 1 of the 7 possible outcomes (see Table C1). or example, in outcome 2, Hat A is red, Hat B is black and Hat C is black. The players under Hat A will see that: Hat B is Black and Hat C is Black. The players under Hat B will see that: Hat A is Red and Hat C is Black. Finally, the players under Hat C will see that: Hat A is Red and Hat B is Black. There is an equal chance for any one of these outcomes. Notice that at each outcome, there will always be at least one black hat. See Figure C1 for an example of what you might observe. Here you are under Hat A and you see that both Hat B and C are black.

Overview of the Round

After you have observed the hat colours, the round will continued to the "trading stage" followed by the "decision stage". You begin the trading stage with 1 Token and a loan of 6000 ECU cash that must be returned at the end of the round. In the trading stage you have the opportunity to either buy more tokens or sell your token. You will only be trading with the other players under the same hat. After all transaction of tokens are completed, only players with at least one token will proceed to the decision stage - if you do not wish to participate in the decision stage, you should sell your token. In the decision stage, you will perform the task of determining your hat colour. After you have completed the decision stage, you will return the loan of 6000 ECU, and your tokens owned will redeemed by the computer (bought by the computer) at a rate that will depend on your behaviours in the decision stage. In the following, we shall first describe the design of the trading and decision stages. Thereafter, we will describe how you token redemption rate will be determined and finally we will describe your payoffs in the round.

Trading Stage

All players begin the trading stage with One Token and a loan 6000 ECU (Money) that must be paid back at the end of the round. Here you are permitted to buy or sell tokens, but only with the other players under the same hat. This implies that the market will consist of exactly 6 players and will last for 120 seconds. You will buy and sell tokens through a continuous double auction mechanism which we will now explain. See figure C2 for a screenshot of the trading stage.

The buy or sell tokens, you will need to first announce you “Ask” and “Bid” prices to all other players. Your Ask price (between 0 and 1200ECU) tells all other players how much you are willing to sell a token for. Your Bid price (between 0 and 1200ECU) tells all other players how much you are willing to buy a token for. The column “Market Ask Prices” reflects the ask prices of all six players you interact with. The column “Market Bid Prices” reflects the bid prices of all six players you interact with. To buy a token, simply select the price on the “Market Ask Prices” column and click “Buy”. Likewise to sell tokens simply select the price on the “Market Bid Prices” column and click “sell”. The column “Market Price” provides the history of all transaction prices for tokens. After 120 seconds, the trading stage will end and you will see on your screens the amount of money you have and the number of tokens you own. See figure C3 for a screenshot.

Decision Stage

Only players with at least one token can participate in the Decision Stage. If you do not have any tokens, you can observe the decision of all other players participating in the Decision Stage through your computer screens but may not yourself participate. You task in the decision stage is to determine the colour of your hat. The decision stage will consist of 4 periods. At each period the computer will present you with the following question, to which you must choose from 3 possible actions (a), (b) or (c).

Computer’s question: *“Do You Know your hat colour?”*

Your actions (a) My Hat is RED, (b) My Hat is BLACK and (c) No! I will decide in a later period Here are some rules:

RULE 1: *At each period, you have a maximum of 4 minutes to choose an action*

RULE 2: *You will immediately end the decision stage if the actions (a) or (b) were chosen*

RULE 3: *You will only go to the next period if you had chosen (c) in the previous period*

RULE 4: *If you arrive at period 4, you can only chose from the actions (a) or (b)*

RULE 5: *Any action chosen will be known to all other players in the subsequent period.*

Here, are some screenshots to help you understand the decision stage design (Figure C4, C5 and C6).

You are said to have “determined your hat colour” when you choose (a) or (b). This is why you will only go to the next period if you have chosen (c). To help you understand the experiment design we have include some screen shoots in Figures C4, C5 and C6.

Figure C4 presents an illustration of the first period in the decision stage. You are under Hat A and you see the other hat colours. In addition the computer presents you with the question “Do you know you hat colour” to which you must reply with one of the 3 possible actions.

Figure C5 presents an illustration of the second period in the decision stage. You are under Hat A and you see the other hat colours. The computer again presents you with the question “Do you know you hat colour” to which you must reply with one of the 3 possible actions. In addition, you also see on your screen the actions chosen by the other players in the previous period. Here there are only two players under hat A who had participated in the decision stage. One player has 4 tokens and the other player has 2 tokens. You see that the player with 4 tokens had chosen (c) in period 1 and the player with 2 tokens had chosen (c) in period 1. Under hat B, there are three players who had participated in the decisions stage. All three player have 2 tokens and had chosen (c) in period 1. Finally, under Hat C, there are 4 players who had participate in the decision stage, one of them has 3 tokens, whilst the other three have only one token. You see that the 3 token player had chosen (c) in period 1. Two of the players with one token had chosen (c) in period 1 whilst the last player, also with one token, had chosen (b) in period 1.

Figure C6 presents an illustration of the third period in the decision stage. You are under Hat A and you see the other hat colours. The computer again presents you with the question "Do you know you hat colour" to which you must reply with one of the 3 possible actions. There are two player under hat A. The two token player had chosen (c) in period 2. The four token player had chosen (b) in period 2. There are 3 players under hat B, each of them with two tokens. One of them had chosen (b) in period 2, whilst the other two had chosen (c) in period 2. There are four players under hat C. The three token player had chosen (c) in period 2. Amongst the one token players, one of them did not participate in period 2 as he had chosen either (a) or (b) in the period 1. Thus that player is said to have ended the game. However, the other two players with one tokens had chosen (b) in period 2.

Token Redemption Rate

After all players have completed the decision stage, your tokens will be redeemed by the computer. The redemption rate will depend on the period which he had "determined your hat colour" and whether you were correct. The payoffs can be easily summarised as followed. Each token is initially worth 950 ECU. The token's value decreases by 50 ECU for each time you had chosen (c). In addition, the tokens value decreases by 700 ECU if you had chosen (a) or (b) and was found to be incorrect, or 0 ECU if found to be correct. See Table C2 for an overview of the redemption rate. Here are some examples to help you understand the redemption rate:

(i) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (b). You have correctly determined your hat colour and your token redemption rate is therefore 850 ECU.

(ii) You hat is black. In period 1 you choose (c), in period 2 you choose (c) and in period 3 you choose (a). You have incorrectly determined your hat colour and your token redemption rate is 150 ECU.

End of Round Payoff

Your payoffs at the end of each round will be determined as followed:

Payoffs = (Money After Trading Stage - 6000) + (Tokens) x (Redemption Rate)

If your payoffs will found to be negative, we will round it off to 0 ECU. This completes the description of each experimental round. After the completion of 10 experiment rounds, we require you to complete a survey before you receive your cash payments. Please feel free to clarify any question or doubts you might have with regards to the instructions.

Table C1: TRADE: The 7 Possible Outcomes

Outcomes	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Chance	1/7	1/7	1/7	1/7	1/7	1/7	1/7
Hat A	B	R	B	B	B	R	R
Hat B	B	B	R	B	R	B	R
Hat C	B	B	B	R	R	R	B

B=BLACK and R=RED

Table C2: TRADE: Token Redemption Rate

	Correct	Incorrect
"Determined Hat Colour" in Period 1	950 ECU	250 ECU
"Determined Hat Colour" in Period 2	900 ECU	200 ECU
"Determined Hat Colour" in Period 3	850 ECU	150 ECU
"Determined Hat Colour" in Period 4	800 ECU	100 ECU

Figure C1: Screen Shot (TRADE) - You see all other Hat colours

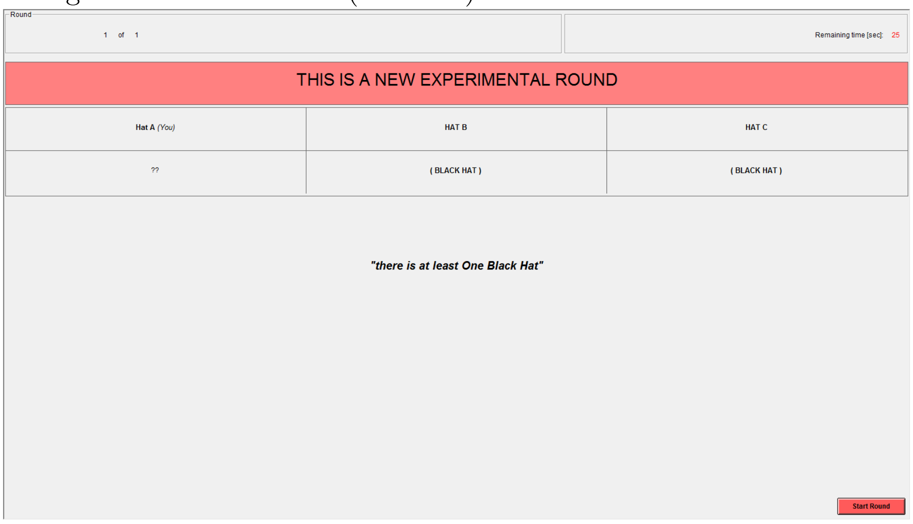

Figure C2: Screen Shot (TRADE) - Trading Stage

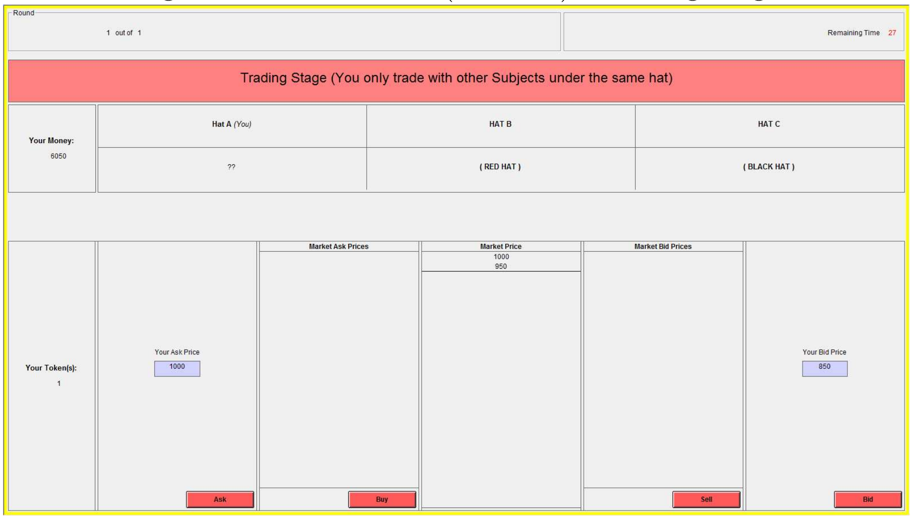

Figure C3: Screen Shot (TRADE) - Trading Stage Results

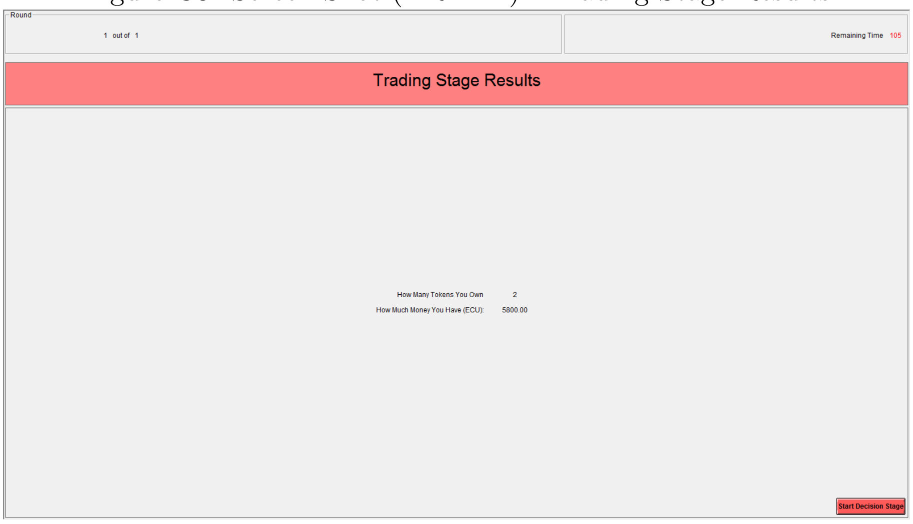

Figure 4: Screen Shot (TRADE) - Decision Stage Period 1

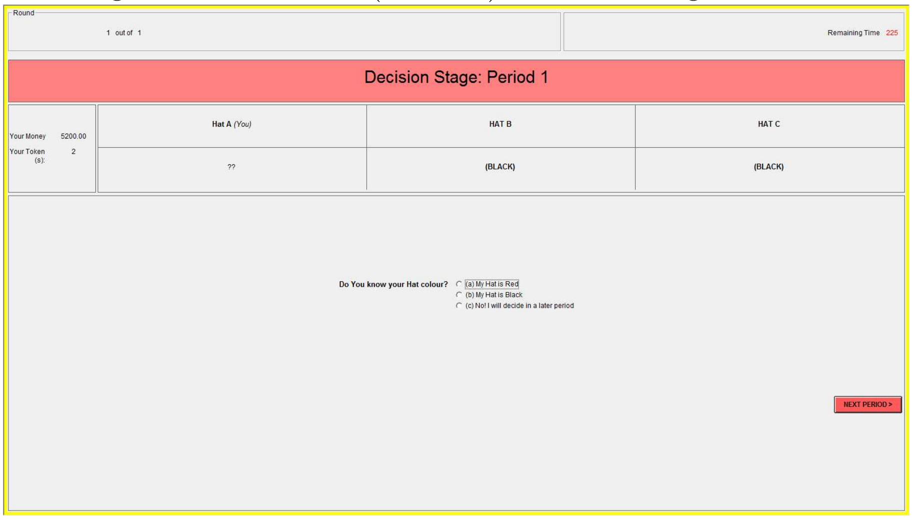

Figure 5: Screen Shot (TRADE) - Decision Stage Period 2

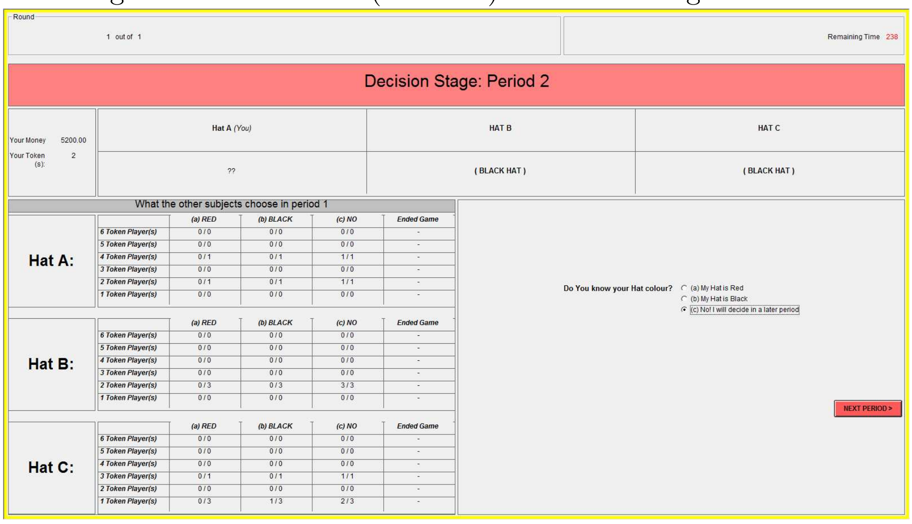

Figure 6: Screen Shot (TRADE) - Decision Stage Period 3