The Pennsylvania State University

The Graduate School

ESSAYS ON BEHAVIORAL MARKET DESIGN AND APPLIED MICROECONOMIC THEORY

A Dissertation in

Economics

by

Robizon Khubulashvili

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

August 2020

The dissertation of Robizon Khubulashvili was reviewed and approved by the following:

Kalyan Chatterjee
Distinguished Professor of Economics and Management Science
Dissertation Adviser
Co-Chair of Committee

Ran Shorrer
Assistant Professor of Economics
Co-Chair of Committee

Karl Schurter
Assistant Professor of Economics

Chloe Tergiman
Assitant Professor of Business

Marc Henry
Professor of Economics
Chair of the Graduate Program

Abstract

In Chapter 1 of my dissertation – co-authored with Ala Avoyan and Giorgi Mekerishvili, we investigate market design for online gaming platforms. Since a large part of such platforms' income is generated by advertisement, it is essential to know what influences how long users stay on the platform. We provide evidence of history-dependent stopping behavior in non-monetary environments. We find that there are two types of people: those who are more likely to stop playing after a loss; and others, who are more likely to keep playing until a win. We find that an individual's type is time-invariant over the years. We propose a behavioral dynamic choice model where utility from playing another game is directly affected by the outcome of the previous game. We structurally estimate this time non-separable preference model and conduct counterfactual analyses to evaluate alternative market designs. We find that in the context of online chess games a matching algorithm that incorporates stopping behavior can increase the length of play by 5.44%.

Chapter 2 of my dissertation, studies how investors react to the arrival of new information when they are considering whether to pull out their investment or not. Proprietary data from online bookmaker shows that bettors' (investors') decision to cash-out their bet (investment) follows systematic patterns. To better understand how new information affects investors' decisions, I create a new game, the "cash-out" game, and study it in the lab. I find that we can divide the population of investors into two main groups. The first group is more likely to withdraw an investment if they get good news — these investors under-interpret good news. The second group is more likely to withdraw an investment after getting negative news — these investors over-interpret bad news. Those biases in belief updating can be used by a market designer to extract larger fraction of consumer surplus by offering dynamic contract which takes into account consumer's type.

In Chapter 3 of my dissertation, I investigate potential collaboration dynamics among two R&D teams called research joint venture (RJV). I showed that with observable types collaboration is sustainable for some parameter values. Next, I consider two cases with asymmetric information, in which one player's ability is observable (player A), while another player's ability is private information (player B). First, I show that if ability of player A is moderately low, RJV and first best outcome is sustainable for limited period of time but there is break up with positive probability in every equilibrium. Second, when ability of player A is very high, RJV is sustainable for some period of time but no matter what player B's actual type is there will be break up and re-establishment of the RJV with positive probability. I also designed an experiment to test theoretical predictions in the laboratory. In line with the theory prediction, I found that reputation building might lead to the break up of the RJV.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the support of wonderful people.

My sincere thanks go to the members of my dissertation committee: Professors Kalyan Chatterjee, Ran Shorrer, Karl Schurter, and Chloe Tergiman. They generously gave their time to help me improve all the chapters of my dissertation. I want to thank, Prof. Chloe Tergiman who encouraged me to keep working. Her support is invaluable for my academic career. Also, I want to thank Prof. Ran Shorrer who patiently helped me during my job market, his support made my job dream reality.

I am very grateful to my friends and co-authors Ala Avoyan and Giorgi Mekerishvili for trusting my idea and building an amazing paper that helped me shape my research. I want to thank Nino, Salome, and their families for welcoming me to their families. I cannot forget my childhood friends, Nika, Leqso, Lasha, Guka, and Goga, for always being there for me.

No words can describe how much my mother - Giuli Giunashvili, father - Archil Khubulashvili and sister Sophio Khubulashvili did for me. Thank you for your selfless sacrifice. Your love and continuous support opened up opportunities for me and let me live the life that I could have only dreamed about. I would like to express my gratitude to my aunt Ani Giunashvili and her family for everything they did for me and my family.

Finally, I want to thank my girlfriend Sophie for always believing in me even when I doubted myself. I am happy that I have my entire life ahead of me to show her how much I love her. I want to dedicate this and all my future work to her, thank you for everything you are doing for me!

Chapter 1: Market Design for Gaming Platforms: Using Stopping Behaviour in Matching Algorithms

What determines our decision of when to stop a given endeavor? Does our past success motivate the stopping decision, or is failure the primary determining factor? In this paper, we use online chess platform data to show history dependence in stopping decision. Furthermore, we provide evidence that different players react differently to wins and losses when they decide whether to play another game. We argue that such history dependence is not explained by existing economic theories, and we propose an alternative theory with time non-separable preferences that rationalizes our data.

We collect the data from the most prominent online chess platform, chess.com, which has over 30 million users and hosts, on average, 3 million chess games every day. We scraped the entire history of play for a random sample of about 20,000 users in the years of 2017 and 2018. Based on the data for 2017, we find that about 60% of players are *loss-stoppers* (players who are substantially more likely to stop playing after a loss), and 25% are *win-stoppers* (players who are considerably more likely to stop playing after a win). The remaining players are *neutral* types. We use the 2018 data to reclassify the same players and find that individuals are, in their vast majority, stable over time in terms of how they react to wins and losses.

We develop a structural model that allows for time non-separable preferences, where *future* game utility can depend on the history of play. Our estimates from the structural model are consistent with the above mentioned reduced form evidence. For some people, loss in a given game decreases the utility from playing another game. For others, it increases utility from playing another game.[1]

We then use our findings to propose a matching algorithm that increases the number of games played. Currently, the platform only uses players' ratings to match similarly rated players with higher probability, ignoring additional information such as a player's "type."[2] The matching mechanism we put forward leverages the following simple observation: loss-stoppers play more when they win, while win-stoppers play more when they lose. Thus, loss-stoppers can be matched with relatively lower rated individuals, so that they have a higher likelihood of winning and a higher likelihood of continuing to play. On the other hand, win-stoppers can be matched with players with higher ratings, so that they are more likely to lose and start one more game. We use our structural model to conduct counterfactual analyses. We show that incorporating user's type can lead to a 5.44% increase in the number

[1] Our classification may best be understood not as a reference to a fundamental, underlying bias, but as a reference to a reduced-form phenomenon. It is worth noting that several different underlying psychological forces could generate this behavior. See the discussion at the end of Section 1.2.1.

[2] The rating system reflects how well a person plays chess.

of games played with only minimally changing the user experience.[3]

Fundamentally, this paper presents and estimates a dynamic discrete choice model in which the agent may have non-separable preferences over the stochastic outcomes of their actions. In that sense, the application is analogous to the optimal stopping problems faced, for example, by taxi drivers, whose decisions to end their shifts may be influenced by their recent fares (see Camerer et al. (1997)). Recent empirical research on this topic is complicated by spatial search frictions and limited by the imperfect observability of the decision-makers' identities and histories of outcome.[4] In contrast, we perfectly observe actions, payoff relevant outcomes, and independent realizations of each agent's decision problem. We take advantage of this rich data to demonstrate that the agent's decisions cannot be reconciled in a model without non-separable preferences and that there is substantial heterogeneity in preferences across players. Although our quantitative findings do not speak to the design of taxi markets, our counterfactual analysis provides an illustrative example of a policy that leverages such preference heterogeneity to further the market designer's goals.[5]

In the context of chess games, Anderson and Green (2018) show that players are more likely to stop playing when they set a new personal best rating. For revenue maximization, these patterns are not useful to design a better market as such events are very rare. Anderson and Green (2018) show that, on average, a person reaches their personal best twice every 15 years.

The rest of the paper is organized as follows. Section 2 provides details on the data collection and presents descriptive results. In Section 3, we introduce the structural model and the identification strategy. In Section 4, we show the results of the structural estimation and counterfactual analysis. In Section 5, we assess the time stability of behavioral types and the validity of our structural modeling choices using the Cox proportional hazard model. Section 6 concludes. We provide further robustness checks in the appendix including the possible effects of rating (ability).

1.1 Data

In this section, we first provide details on the data collection and information about the platform. We then provide descriptive statistics that illustrate consistent behavioral patterns.

[3]User behavior may change if we make changes to the matching algorithm, especially if the changes are drastic and can be inferred by the users. In this paper, we abstract from this concern, but we still focus on alternative matching algorithms that we believe would minimally affect the user experience.

[4]See Thakral and Tô (2017); Farber (2005, 2008, 2015); Crawford and Meng (2011); Abeler et al. (2011); Morgul and Ozbay (2015); Cerulli-Harms et al. (2019); Frechette et al. (2019).

[5]In principle ride-sharing platforms such as Uber and Lyft can use a similar approach to increase drivers labor supply, for example by using the information on user's tipping behavior to match riders with drivers.

1.1.1 Collection

We scraped the data from an online chess platform, chess.com, which is the most frequently visited chess website.[6] The website has over 30 million users, and it hosts around 3 million chess games every day. Users range from amateur players to the world's best chess players, including Magnus Carlsen, the World Chess Champion.[7] This platform is free to use, and anyone can register to play against other people or a computerized opponent. The website also provides some lessons and chess puzzles.

We use the public Application Programming Interface (API) to collect data. Each observation includes information about the players and the game: usernames, their self-identified country of association, platform ratings, the time at which the game was played, which player had white pieces, the length of the game, and the final results. We focus on a sub-sample of the data that contains "blitz"[8] games and games played against human players (see Table 1 for a summary).

1.1.2 Definitions

Number of games	35,517,526
Number of sessions	6,063,366
Number of players	13,027
Average Session Length	4.14
Average Number of Session	348
Average Rating	1311
Rating Range	[100, 2816]
Pr(Win | White Pieces)	0.51
Pr(Win | Black Pieces)	0.47
Percentage of draws	3.4

Table 1: Data description

A *game* g is a single game with a human opponent. A collection of games ordered by time stamp, $(g_1, g_2, ..., g_n)$, is called a *session* if there is no game played T minutes before g_1 or after g_n, and for any $i \in \{1, ..., n-1\}$, the time between g_i and g_{i+1} is less than T.[9] We call sessions that contain only one game ($n = 1$) *only game* (O-game). For sessions with

[6]Based on Alexa internet rating, www.alexa.com.

[7]Magnus Carlsen, a Norwegian chess grandmaster, and chess prodigy is the highest-rated player in the world, and the highest-rated player in the history of chess.

[8]Blitz is a type of chess game where each player has somewhere between 3 to 15 minutes for the game.

[9]For the main section of the results we set $T = 30$ minutes, we vary T to check the robustness of our results and we find no substantial differences.

$n \geq 2$, g_1 is the first game, g_n is the last game and any game in between the first and the last is referred to as *a middle game* ($k \in \{2, 3, .., n-1\}$).

Based on the terms above, we categorize sessions into three mutually exclusive groups: (1) sessions that consist of a single game – the only game $S = (O)$; (2) session with exactly two games $S = (F, L)$; (3) sessions with more than two games $S = (F, M_1, ..., M_m, L)$ where $m \in \{1, ..., n-1\}$.

Let $f_W(\cdot)$ be a function that calculates the fraction of wins in a particular type of game, for example, $f_W(L)$ is a player's fraction of wins in the last games.[10] In some cases, when the context is clear, instead of writing $f_W(F)$, $f_W(M)$, $f_W(L)$, $f_W(O)$, we write F, M, L, and O, and we mean the fraction of wins in first, middle, last, and only games, respectively.

1.1.3 Descriptive Results

In this section, we first establish that session-stopping behavior is history-dependent. We find that a sizable fraction of players fall into two categories: some people consistently leave their session after they win a game, while others exit after a loss; we refer to these types as *win-* and *loss-stoppers*, respectively. In section 1.2, we present a theoretical model that can account for the behavior observed in the data.

History Dependence

We first take all sessions in our data that lasted at least three games and calculate the average winning frequency in the first, middle, and last games for *each player*, as defined in Section 1.1.2. If a player makes a decision to stop the game randomly and stopping behavior is history independent, then the fraction of wins in the last game should be similar to the fraction of wins in any other game. Our null hypothesis is:

H_0: Players stopping behavior is independent of the outcome of the last game.

H_0 implies that the correlation between the fraction of wins in the last and middle games to be close to 1.[11] Figure 1a presents the relationship between the fraction of wins in the last and middle scores (the solid line represents the linear regression line, where the dependent variable is the fraction of wins in the middle games and the independent variable is the fraction of wins in the last game). The correlation between the fraction of wins in the last and middle games is $-.41$ and statistically different from 1 with $p < 0.001$. Thus, at the aggregate level, the decision to stop is not random and H_0 is rejected.

[10] For the most part of the paper, we look at a user's playing history in year 2017. For each user we calculate the fraction of wins in different types of games.

[11] We do not include first games at this stage because the data on first games is used later to check the robustness of the findings.

(a) Fraction of wins (b) Density of fraction of wins

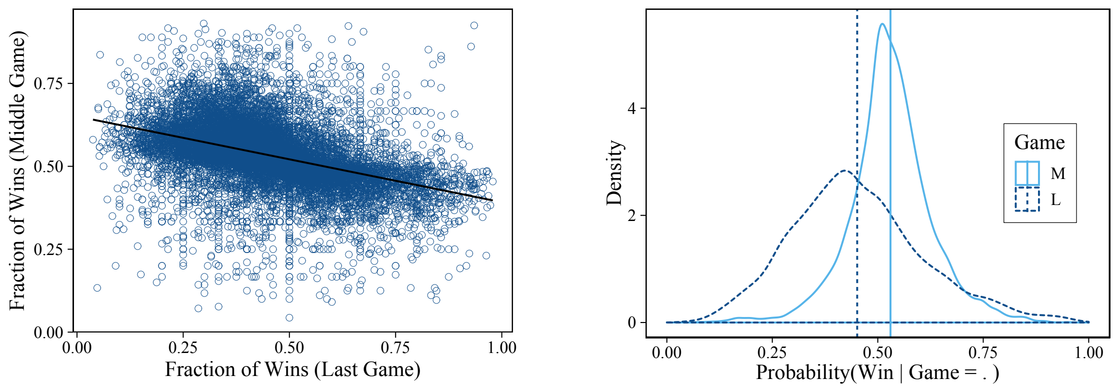

Figure 1: Fraction of wins in the Last and Middle games.

Behavioral Types

We rejected the null hypothesis of history independence; however, the alternative hypothesis does not identify the nature of the relationship between the outcome of the game just played and the decision to play another game. Are players more likely to end the session after a win or after a loss? We state the following two alternative hypotheses:

H_A^1: *players are more likely to end a session after a win.*

H_A^2: *players are more likely to end a session after a loss.*

If only one of the two alternative hypotheses is correct, we should see a skew or a shift of the distribution of the fraction of wins in the last game compared to the fraction of wins in the middle games (Figure 1b presents both distributions). The test of medians for the two distributions shows that the median fraction of wins in the last game is statistically lower than the median fraction of wins in the middle game. If we only look at this aggregate result, it supports the hypothesis that players are more likely to stop playing after a loss. However, if we take a closer look at the distribution, we see an interesting pattern. The standard deviation of the fraction of wins in the last game is twice that of the middle game. Both plots in Figure 1 inform us that there are some people with a much higher fraction of wins in the last game compared to the middle game and that there are others with a much lower fraction of wins in the last game compared to the middle game. We define players that are more likely to end the session on a loss and others who are more likely to end the session on a win in the following way:

A player is behavioral type at the tolerance level of τ and he is

- a win-stopper if $f_W(L) > f_W(M) + \tau$,
- a loss-stopper if $f_W(L) < f_W(M) - \tau$,

neutral types have $f_W(L) \in [f_W(M) - \tau, f_W(M) + \tau]$.

Using data from sessions that last 3 or more games, we classify players according to Definition 1.1.3 (see Figure 31 in the appendix for population decomposition by types using different tolerance levels of τ). Intuitively, as we increase the tolerance level, fewer players are classified as win- or loss-stopper types. Interestingly, when we change τ the ratio of win-stoppers to loss stoppers stays stable at around 40%. At tolerance level τ of 5%, 84% of the players are classified as behavioral types, with about 30% of them being win-stoppers and a larger fraction, 70%, loss-stoppers.[12]

To further examine that Definition 1.1.3 captures the patterns in the data, let us examine only-games—sessions that contain only one game. Note that we have not used the sessions with only one game for the classification; so far, we used sessions of length three and above. For ease of exposition, let us take Definition 1.1.3 to the extreme, where we assume that a win-stopper type **always** stops after a win and a loss-stopper type **always** stops after a loss. This extreme definition of types implies the following two observations (see Figure 2). First, the fraction of wins for win-stoppers in the only and last games must be 1. This is because if a player wins this game, he ends a session, and the game is classified as the only game. If a player loses the first game, he will start another game, making this session at least two games long, in which case whenever a player wins, we classify that game as the last game. Therefore the fraction of wins in the last game is also 1. Second, the fraction of wins for loss-stopper types in the only and last games must be 0 (the argument is similar to the win-stopper case). Combining these two observations leads to the following prediction.

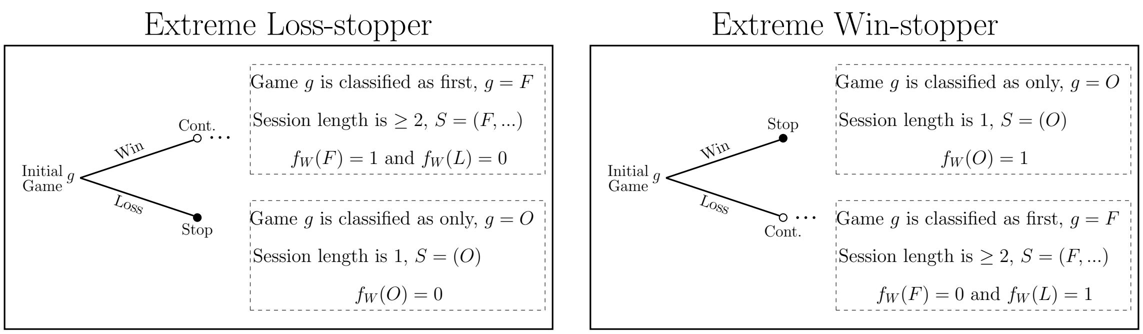

Figure 2: Extreme Behavioral Types' Actions

The correlation between the fraction of wins in the last and only game is positive.

We find that win-stoppers' fraction of wins in the only-game is two times higher than that of loss-stoppers. Figure 3 presents a scatter plot with the fraction of wins in the last-game and the fraction of wins in the only-game. A strong and significant positive relationship

[12] Unless specified otherwise, we use 5% tolerance level.

implies that the types who are more likely to stop playing on a win similarly have a higher fraction of wins in the only game and vice versa, as stated in Prediction 1.1.3.

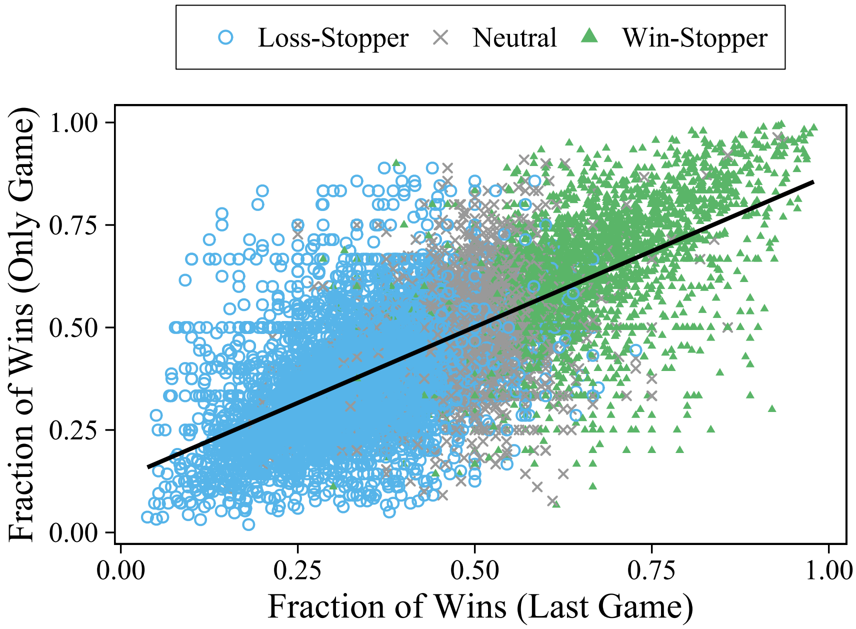

Figure 3: Fraction of wins in the Last and Only Game

So far, we used a fraction of wins in the middle, last, and only games. We have not used a fraction of wins in the first games. The diagram in Figure 2 describes what each type would do if they won or lost the initial game. It gives a relationship between the fraction of wins in the first and only games. If the loss-stopper wins the initial game he plays another one, which makes that game classified as the first game, while if loss-stopper loses the initial game, he stops playing, which makes that game classified as an only game. Therefore loss-stopper's fraction of wins on the first game must be one, and in the only games must be 0. Similarly, we argue that for win-stoppers, the fraction of wins in the first game must be 0, while the fraction of wins in the only game must be 1. Note that while the prediction for last- and only-game relationship is positive, the prediction for the first- and only-game relationship is negative.

The correlation between the fraction of wins in the first and only game is negative.

Figure 4 presents the correlation matrix with p-values in parentheses. A negative relationship between the fraction of wins in the first and only games implies that the types who are more likely to stop playing on a win have a higher fraction of wins in the only-games since they would have kept playing if they had lost the initial game. Similarly, the types who are more likely to stop playing on a loss have a lower fraction of wins in the only-games.

Similar to predictions 1.1.3 and 1.1.3, we can use the diagram in Figure 2 to write four

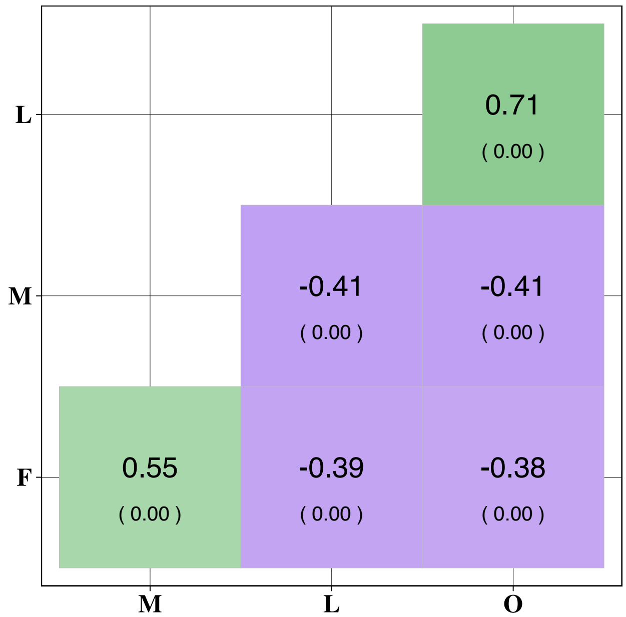

Figure 4: Correlation Matrix

more predictions about relationships between the fraction of wins in the first, middle, last, and only games. All those predictions are accurately matched to the data (see Figure 4).

1.2 The Model

We do not take a stand on whether people end a session after a win or loss. Instead, we model different types of behavior and let the data inform us about the types. We capture the different types of people by allowing the utility from playing a new game to depend on the last game outcome and the person's type. We call a person *loss-stopper* type, if the utility from playing one more game for that person decreases after a loss compared to that after a win. Similarly, we call a person *win-stopper* type, if the utility from playing another game increases after a loss compared to that after a win. We say that a person is *neutral* type if the person's utility is not affected by the outcome of the last game.

1.2.1 Description

This section lays out a chess player's dynamic choice problem. A chess player is characterized by his type that consists of an element observable to the player, player's opponents and econometrician, and an element privately known to the player.

Let y be a vector of characteristics that may change over time, and are observable to the player, opponents, and econometrician (e.g., player's rating). We assume that y lives in finite space Y. A player can be one of the following 3 types: win-stopper (θ_W), loss-stopper (θ_L), or neutral (θ_N). Let $\Theta = \{\theta_W, \theta_L, \theta_N\}$ be the set of all types and let θ be an element of this set. The player's type is fixed over time. A player's type profile at time t, (y_t, θ), consists of player's time-variable characteristics, y_t, and fixed unobservable type, θ. We use variables without time subscripts to denote current states and 'prime' superscripts to denote the next period's state.

Each period, a player is facing the following decision: given the previous history of the play, the player needs to decide whether to play an additional game or to go offline and take an outside option. Before making this decision, the player's utility from another game is

$$U(y, \theta, \chi) = u(y) + (1 - \chi)l_\theta \tag{1}$$

where y is the player's current rating, θ is the player's type, and χ is the outcome of the last game. If a player won the last game ($\chi = 1$), the utility from playing another game is $u(y)$. We can think of this term as how much the player likes playing chess independently of his type. If the player lost the last game, then his utility from playing another game depends on his type.

A player is

i) Loss-Stopper, if $l_\theta < 0$;

ii) Win-Stopper, if $l_\theta > 0$;

iii) Neutral, if $l_\theta = 0$.

There is an outside option, c, that every period is independently drawn from a distribution with density $f(c)$. If a player ends a session, he takes an outside option c. Otherwise, the player's utility is $U(y, \theta, \chi)$ from playing a new game and he moves to the next period, at which point the player faces the same decision based on the new history of the last game (χ'). In each period (after a game is over) a player is facing the following problem:

$$V(y, \theta, \chi, c) = \max\Big\{c,\ u(y) + (1 - \chi)l_\theta + \delta \sum_{\substack{\chi', y' \in \\ \{0,1\} \times Y}} p(y'|y, \chi')p(\chi'|y)V(y', \theta, \chi')\Big\} \tag{2}$$

where δ is the discount factor; $p(y'|y, \chi')$ is the probability of receiving a rating y' given that the player's current rating is y and the outcome of the next game is χ'; $p(\chi'|y)$ is the probability of the subsequent game outcome given the current period rating y. Note that the player matching mechanism, which is based on the player's rating, affects the probability of winning the next game. This is the way through which changing the matching mechanism can influence a player's decision to start a new game. Note that we can directly recover the law of motion of y from the data (player's rating updating rule).

It is important to note that we focus on a reduced-form phenomenon: win-stoppers are more likely to stop playing after a win compared to a loss, and loss-stoppers are more likely to end the session after a loss compared to a win. Some papers evaluated what psychological biases may lead to behavior similar to loss-stopper or wins-stopper separately but not the heterogeneity we observe in this paper (see section 3.6 in the appendix for a theoretical discussion and more details).[13]

1.2.2 Identification

In this section, we provide the identification of players' types, l_θ, δ, probabilities of winning, outside option distribution parameter, and matching probabilities. We start with the identification of behavioral types.[14]

Optimal stopping rule is a threshold rule in c.

Proof. Note that in equation (2), continuation values do not depend on the **current** realization of c. Hence, fixing the continuation values and current period utility from playing another game, the second term under the max operator is lower than outside option c, for sufficiently high c. So, we have a threshold, $\bar{c}(y, \theta, \chi)$, such that for realizations of c above this threshold, the player stops playing and takes the outside option. ■

Therefore $\bar{c}(y, \theta, \chi)$ is a threshold such that a player with type profile (y, θ) having an outcome χ in the last game ends a session if and only if the realized c is at least as large as $\bar{c}(y, \theta, \chi)$.

From (2), we have,

$$\bar{c}(y, \theta, \chi) = u(y) + (1-\chi)l_\theta + \delta \sum_{\chi', y' \in \{0,1\} \times Y} p(y'|y, \chi')p(\chi'|y)V(\theta, y', \chi') \tag{3}$$

The following proposition leads to the identification of behavioral types.

[13]See Miller and Sanjurjo (2018); Miller et al. (2018); Miller and Sanjurjo (2017); Aharoni and Sarig (2012); Arkes (2010, 2013); Avugos et al. (2013); Cervone et al. (2014); Brown and Sauer (1993); Camerer (1989); Croson and Sundali (2005); Suetens et al. (2016); Gilovich et al. (1985); Green and Zwiebel (2017); Koehler and Conley (2003); Rabin and Vayanos (2010); Rao (2009b,c,a); Rinott and Bar-Hillel (2015); Sinkey and Logan (2014); Stone (2012); Stone and Arkes (2018); Sundali and Croson (2006); Tversky and Gilovich (1989); Wardrop (1999); Xu and Harvey (2014); Yaari and Eisenmann (2011).

[14]The identification and estimation of the theoretical model are in the tradition of Hotz and Miller (1993). We show how we can forgo numerical dynamic programming to compute the value functions for every parameter vector, and propose an estimation procedure that is simple to implement and computationally efficient.

Proposition 1
i) $\bar{c}(\theta_W, y, 0) > \bar{c}(\theta_W, y, 1)$*;*
ii) $\bar{c}(\theta_L, y, 0) < \bar{c}(\theta_L, y, 1)$*;*
iii) $\bar{c}(\theta_N, y, 0) = \bar{c}(\theta_N, y, 1)$*.*

Proof. The proof follows from equation (3). ∎

Proposition 1 implies that the win-stopper types' probability of playing one more game is higher if they lose the previous game compared to when they win, and vice versa for the loss-stopper types. For the neutral types, that probability is the same no matter the history of outcomes. By Proposition 1 we can identify a behavioral type of a player from the data by looking at his stopping probabilities after losses and wins.

Model Parameters

The following parametric assumption is made on the distribution of outside option, $F(c)$, $F(c)$ is an exponential distribution with parameter λ.

We now argue that under the assumption 1.2.2 and by normalizing one parameter of our choice in the model, we can identify δ, λ, l_θ and $u(\cdot)$. In the estimation, we normalize $\lambda = 1$. Let,

$$H(y, \theta) = u(y) + \delta \sum_{\chi', y' \in \{0,1\} \times Y} p(y'|y, \chi')p(\chi'|y)V(\theta, y', \chi') \tag{4}$$

Under the assumption 1.2.2 and from equation 3, we get that the probability of stopping and taking outside option, $h(y, \theta, \chi)$, can be written as

$$h(y, \theta, \chi) = e^{-\lambda(H(y,\theta)+(1-\chi)l_\theta)} \tag{5}$$

$\lambda H(y, \theta)$ and λl_θ are identified for all (y, θ).

Proof. Let us look at equation (5) for $\chi = 1$. LHS, $h(\theta, y, 1)$ is coming from the data as probability of stopping after a win. On the RHS, the second term in the power, $(1-\chi)l_\theta = 0$, therefore we can say that $\lambda H(\theta, y)$ is identified from (5).

Now substitute $\chi = 0$ in equation (5). Again the LHS is coming from the data as probability of stopping after a loss $h(\theta, y, 0)$. On the RHS, first term in the power $H(\theta, y)$ is coming from the first part of this prove. Thus, λl_θ is identified from (5) as well. ∎

$\lambda V(y, \theta, \chi)$ are identified for all (y, θ, χ).

Proof. Let us denote $\bar{c}(\theta, y, \chi) \equiv \bar{c}$. We can rewrite (2) as:

$$V(\theta, y, c, \chi) = \mathbf{1}(c > \bar{c}) * c + \mathbf{1}(c \leq \bar{c})\Big(H(\theta, y) + (1-\chi)l_\theta\Big) \tag{6}$$

Where $\mathbf{1}(\cdot)$ is an indicator function. Taking expectations of both hand sides of (6) with respect to c, gives

$$\begin{aligned} V(y,\theta,\chi) &= E(c|c>\bar{c}) + Pr(c\leq\bar{c})\Big(H(\theta,y)+(1-\chi)l_\theta\Big) \\ &= Pr(c>\bar{c})\Big(E(c)+\bar{c}\Big) + Pr(c\leq\bar{c})\Big(H(\theta,y)+(1-\chi)l_\theta\Big) \\ &= Pr(c>\bar{c})\Big(\frac{1}{\lambda}+\bar{c}\Big) + Pr(c\leq\bar{c})\Big(H(\theta,y)+(1-\chi)l_\theta\Big) \\ &= Pr(c>\bar{c})\Big(\frac{1}{\lambda}+\bar{c}-H(\theta,y)-(1-\chi)l_\theta\Big) + H(\theta,y)+(1-\chi)l_\theta \end{aligned} \tag{7}$$

multiplying both hand sides by λ and substituting $\bar{c}(y,\theta,\chi)$ from (3), we get,[15]

$$\lambda V(y,\theta,\chi) = e^{-\lambda[H(y,\theta)+(1-\chi)l_\theta]} + \lambda H(y,\theta) + (1-\chi)\lambda l_\theta \tag{8}$$

Claim 2 and expression (8) imply that $\lambda V(y,\theta,\chi)$ are identified for all (y,θ,χ). ■

δ and $\lambda u(y)$ are identified.

Proof.

We can consider the difference $\lambda(H(y,\theta)-H(\theta',y))$ for some $\theta\neq\theta'$. This gives us,

$$\delta = \frac{\lambda(H(y,\theta)-H(\theta',y))}{\lambda(\sum_{\chi',y'\in\{0,1\}\times Y} p(y'|y,\chi')p(\chi'|y)(V(\theta,y',\chi')-V(\theta',y',\chi')))}$$

By claims 2 and 3, numerator and denominator are identified in the above equation.

We identify $\lambda u(y)$ from,

$$\lambda u(y) = \lambda H(y,\theta) - \lambda\delta \sum_{\chi',y'\in\{0,1\}\times Y} p(y'|y,\chi')p(\chi',y)V(\theta,y',\chi')$$

■

Finally, we can normalize all the parameters and value functions by λ. This completes the identification of the parameters of the model.

1.3 Structural Estimation and Counterfactual Analysis

In this section, we first introduce the structural estimation results. Then we present the counterfactual analysis results that highlight how a market designer can change how much time a player spends on the platform. We compare the types identified by the model to the initial classification from Section 1.1.3 and then show the time stability of the model types.

[15] Notice that $Pr(c\leq\bar{c}) = 1-Pr(c>\bar{c}) = 1-e^{-\lambda[H(y,\theta)+(1-\chi)l_\theta]}$

To better understand the counterfactual analysis, we begin by describing the rating system on the platform. Once a player signs up for chess.com, he gets an initial rating (1200 points). Rating changes after every game based on the outcome of the game. Intuitively, the rating goes up after a win and goes down after a loss. Thus, a player's current rating reflects his current expertise in chess: higher the rating, better the player. We recover the rating updating rules from the data. To estimate parameters structurally, we divide the rating range into the grids of 20 since the rating has a wide range ($[100, 2798]$).[16] Since we have very few observations where the rating is below 800 or above 2200, we place all the players with a rating below 800 in the first rating grid and those above 2200 in the last rating grid (grid 71). The rest of the rating range is divided into 20 point intervals.

1.3.1 Structural Estimates

The main parameters that we estimate are l_θ for $\theta \in \{\theta_W,\ \theta_L,\ \theta_N\}$. The estimates are presented in Table 2. We sub-sample the data from the entire sample 300 times with replacement and find that our estimates are stable.

Parameter	Mean	SD	[min, max]
l_{θ_W}	0.810	0.009	$[0.776, 0.838]$
l_{θ_N}	-0.067	0.004	$[-0.077, 0.057]$
l_{θ_L}	-0.698	0.004	$[-0.706, -0.686]$

Table 2: Bootstrapped values for l_θ

Table 2 shows that for a win-stopper type, utility from playing another game increases by 0.810 after a loss compared to a win. The effect is opposite for loss-stopper types, loss in the last game decreases utility from playing another game by 0.696 compared to a win. Intuitively for neutral types, last game result has only a minor effect on utility.

1.3.2 Counterfactual Analysis

Can the market designer leverage information on behavioral types to increase a player's expected time spent on the platform? To answer this question, we need to know what can be controlled by the market designer. In our setting, the player-to-player matching algorithm is controlled by the market designer. Therefore we need to know what is the current matching algorithm, which can be recovered directly from the data. The platform has a simple matching rule. Two players with closer ratings are matched into pairs with

[16] For the counterfactual analysis, we change the grids, more details are provided in the counterfactual section. Changing the grid size or the number of grids does not qualitatively affect any of our results.

higher probability. In other words, the platform only uses the players' rating to decide who plays with whom.[17] In our counterfactual exercise, we allow the platform to choose from matching mechanisms that can be contingent on players' behavioral type in addition to the rating.

Before presenting counterfactual results, we need to explain the data that we use for counterfactual analysis. In the last subsection, we presented estimates for the entire data, but for counterfactual analyses, we focus on a sub-sample. The reason is that we want players to be homogeneous in their rating. In other words, we want to consider players who have similar rating to avoid matching players, one with a high rating and other with low rating. Therefore, we consider only players who have an average yearly rating from 900 to 1300. Then we divide rating space [800, 1400] into 102 grids, each grid size is 6 points.

We re-estimated the model and find that even in the restricted data, the parameter estimates of l_θ are similar to what we had for the entire data (Table 3). To check the stability, we bootstrap the data 300 times.

Parameter	Mean	SD	[min, max]
l_{θ_W}	0.718	0.005	[0.711, 0.726]
l_{θ_N}	-0.047	0.002	[-0.050, 0.045]
l_{θ_L}	-0.613	0.002	[-0.615, -0.609]

Table 3: Bootstrapped values for l_θ with restricted data

The goal of this section is to provide a matching algorithm that is better than the current one. Ideally, we would find an optimal matching algorithm among the class of algorithms that accounts for rating as well as players' types. However, due to the high dimensionality of the problem (we have $3n(n-1)$ variables to optimize, where n is the number of rating grids), our current computational capabilities do not allow us to find the best such algorithm. Instead, we answer the following question: at least, by how much can the platform increase the expected time spent on the platform? To answer this question, we consider optimization over a smaller class of matching algorithms and find that the platform can increase the time spent by the user theoretically by at least 62%.

There is a concern with this algorithm. The algorithm matches players with very low ratings to players with a very high rating. Such drastic change is matching algorithm could lead players to learn about the matching or become uninterested in a platform that frequently matches them with either very low or very high rating players. Therefore we consider the matching algorithm, which does not allow for such radical matching pairs. Intuitively, such an

[17] Obviously we assume that both players who are matched are on the platform and requested a particular type of chess game, for example, 3 minute blitz game.

algorithm results in less improvement – a 5.44% increase. To put this number in perspective, let us look at simple statistics. An average player plays 348 sessions a year, and an average session lasts about 4.14 games. Since we consider only blitz games and average blitz games in our sample lasts for 8 minutes and 12 seconds, 5.44% improvement results in 87 more games per year (on average 11 hours and 54 minutes more time spend on the platform per year per player).

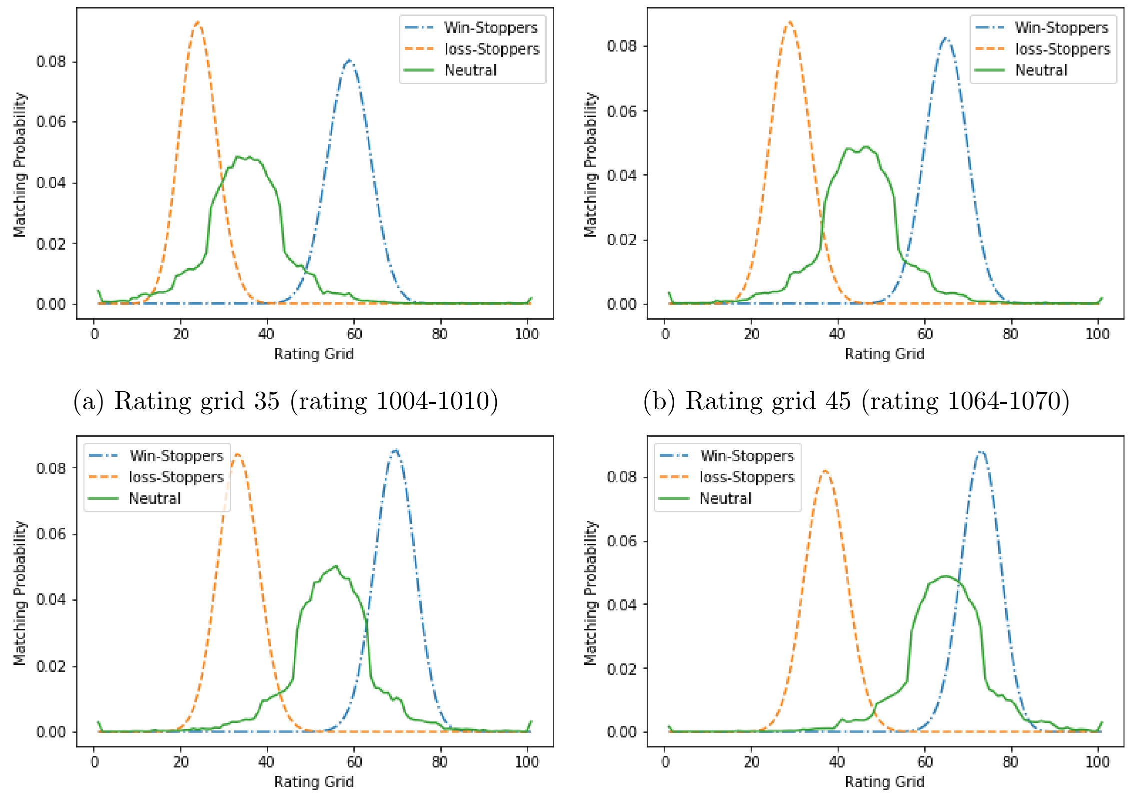

(a) Rating grid 35 (rating 1004-1010)

(b) Rating grid 45 (rating 1064-1070)

(c) Rating grid 55 (rating 1124-1130)

(d) Rating grid 65 (rating 1184-1190)

Figure 5: Matching probabilities for four different rating grids

Figure 5 presents the resulting matching probabilities for four different rating grids. Green (solid) distribution represents current matching probabilities, while orange (dashed) and blue (dotted-dashed) distributions are new (improved) matching probability distributions respectively for loss-stopper and win-stopper types.

We choose middle grids because, for them, we have both higher as well as lower-rated possible opponents. If we match win-stoppers with relatively higher rated players, we decrease the probability that the current player wins leading them to play another game with higher probability. For loss-stoppers argument is reversed. As figure 5 shows, loss-stoppers should be matched with lower-rated players so that they have a higher chance of winning and continue playing with higher probability. Since the neutral types are not affected by the outcome of the last game result, we use these players to clear the market. The matching algorithm explained above and presented in figure 5 improves average session length by 5.44%.

1.4 Robustness of the Model and Time Stability

1.4.1 Time Stability of Behavioral Types

We use data for 2017 for the estimation. However, we also have data for the same people for 2018, which we keep as a holdout sample to check the time stability of behavioral types. Figure 6 shows the relationship between these two years. First, we created individual variable, break probability difference, as the difference between the fraction of times a player took a break after a win and after a loss. Figure 6a shows that there is a positive correlation between those two years. People who have a higher fraction of breaks after a win then after a loss in 2017, also have very similar break probability difference in 2018.

(a) Time Invariance of Types (b) Transition From 2017 to 2018

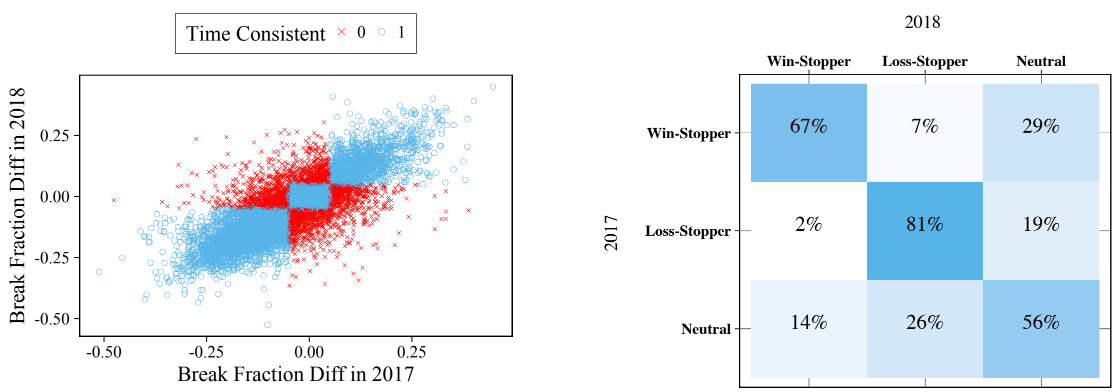

Figure 6: Time Invariance of Types (2017 to 2018)

As Figure 6b shows, the major differences in classification occur for neutral types. This result is not surprising since the definition of types is based on the threshold level. Among neutral types, players who change the type are threshold cases (see Figure 6a).

1.4.2 Consistency of Behavioral Types

In this paper we have two distinct definitions of behavioral types (Definitions 1.1.3 and 1.2.1). The two definitions are intuitively related, but these two methods of identifying behavioral types do not have to overlap at all. That is, given some data, a player could be identified as win-stopper according to the model but be classified as loss-stopper according to the last-game and middle-game definition. For example, suppose a player's complete playing history contains the following set of 3 sessions:

$$\{WWWW, WLW, WLL\},$$

Let us first calculate the stopping probability after a loss, $\Pr(Stop|Loss) = 1/3$. Then, given all the wins and we calculate the stopping probability after a win, $\Pr(Stop|Win) = 2/7$. Given that the probability to stop playing is higher after a loss than after a win, $\Pr(Stop|Loss) > \Pr(Stop|Win)$, our model would qualify the player as loss-stopper. However, according to our definition through the last game of the session, the player's behavior corresponds to win-stopper, because the last game is won more often than middle game.

This example demonstrates that two definitions are intuitively related; however, one does not imply the other. This fact strengthens any relationship we find between the two classifications, highlighting the consistency of our intuition with the proposed theory. Let us look at how the two classifications pair. We take all the players, and in Figure 7a we put a plus sign if the model identifies a win-stopper player as win-stopper, loss-stopper as loss-stopper and neutral as neutral. For 81% of the players in the data, the two classifications match. This result provides strong evidence that the model captures the players' behavior and that the game outcome affects the utility of the next game.

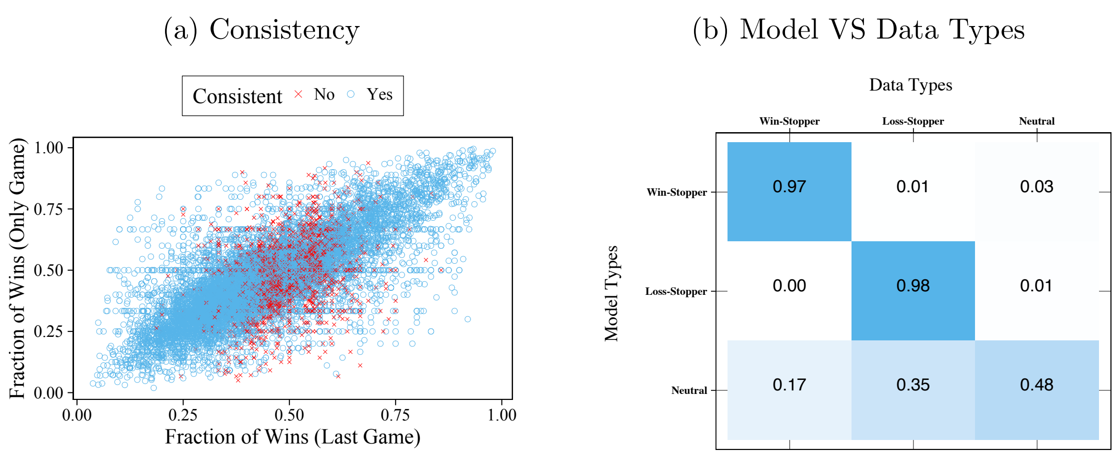

Figure 7: Model Identified Types and Consistency

Figure 7b presents a transition matrix between model types to behavioral types. We observe a large mass on the diagonal, meaning that the two classifications are consistent. For example, 96% of win-stopper types identified by the model were identified as win-stopper types by Definition 1.1.3. However, there are mismatches; for example, some neutral types are classified as behavioral types by the model and vice versa. Notice that the cases in which a win-stopper (loss-stopper) is identified as loss-stopper (win-stopper) by the model account happen only 1% of the time.

1.4.3 Validity of the Modeling Choices

To show that the relevant factor that affect players' stopping decision is the last game result, we estimate the Cox Proportional Hazard (CPH) model. The purpose is to evaluate the effect of a number of factors on survival.[18] Survival analysis allows us to examine how specified factors influence the session stopping rate. Such a rate is referred to as hazard rate, and the examined factors are called covariates.

We use Cox's proportional hazard model with time-dependent covariates. We proceed by first estimating the model using aggregate data and then, later on, we estimate the model for each player separately.[19] The general description of the model is as follows:

$$h_j(t, x_j(t)) = h_0(t) \exp\{x_j(t)'\beta\} \tag{9}$$

LHS of the equation 9 represents risk that game j with characteristics $x_j(t)$ is the last game of the session (session terminates after that games). RHS is comprised of two components: baseline risk and relative risk. The baseline risk, $h_0(t)$, represents the risk that a game will be the last game in a session when all the covariates equal to zero, $x_j(t) = \mathbf{0}$. The relative risk, $\exp\{x_j(t)'\beta\}$, is a proportionate increase or reduction in risk associated with the set of characteristics $x_j(t)$.

Let us look at the CPH results when we pool the data and examine the effect of the last game on the decision to stop a session. The estimation results are presented in Table 4.

Covariate	Coef	exp(Coef)	p-value
Score	-0.23	0.80	< 0.005

Table 4: CPH without type heterogeneity

The variable *Score* takes value 1 if a player wins the last game and 0 otherwise. Interpreting the result of the analysis is easier using the third column (exp(Coef)), where exp(Coef) = 1 implies that whether the last game is a win or a lose has no effect on the decision to stop the session. The value of exp(Coef) is .8, and it shows that a player is 20% less likely to stop playing after a win than after a loss. Recall that this behavior is observed

[18] Survival in our setting means "not ending" a session

[19] Appendix 3.6 shows the results of individual estimates for the CPH model. We also explain how the type identification using the CPH model coincides with our structural model type identification.

among loss-stopper types, but not among the win-stopper types who are more likely to continue playing after a loss. This result is hiding an important heterogeneity that we are aware of based on the analysis in the previous sections of the paper. The estimation does not take into account any differences between players and assumes there is a universal effect of the last game outcome for all players. The negative relationship that we observe is due to the fact that more player are loss-stoppers than win-stoppers.

We introduce behavioral type heterogeneity in players according to our model estimation in Section 1.2. For the ease of exposition of the results, we assume that there are no neutral types, and we have only two types of players: loss-stoppers and win-stoppers. We include the player's type and interaction of type and the last game result to the covariates and re-estimated the CPH model. The results are presented in Table 5.

Covariate	Coef	exp(Coef)	p value
Score	-0.54	0.58	< 0.005
Type	-0.55	0.58	< 0.005
Type $\times$ Score	1.01	2.76	< 0.005

Table 5: CPH with type heterogeneity

The variable *Type* is 1 for win-stopper types and it is 0 for loss-stoppers. Therefore, the baseline in the estimation is loss-stopper losing in the last game. Table 5 shows that for a win-stopper type (Type= 1), the hazard rate is higher after a win than after a loss.[20] In other words, a chance that a win-stopper type ends a session after a loss is lower than after a win. That is what we expected. For loss-stopper types, the relationship is reversed. A win in the last game decreases the hazard rate compared to a loss (baseline) by 42%.

History dependence of length 2

Up until now we have focused on the effect of win/loss history in one game. Now, we examine whether there is an effect of more than one lag on the decision to stop a session (last two games). We estimated the CPH model as before, but now we add Lag 1 score and interaction of this variable and player's type. Table 6 presents the results of the estimation.

Adding the outcome of the game previous to the last game has little effect on the estimates. The last game outcomes have effects similar to what we observed in Table 5. The lagged values have statistically significant effects, but the magnitude of the effect is much smaller than the effect of the last game outcome. Note that the sign of the effect of the lagged variable is predicted using the arguments presented in Section 1.1.3.

Starting rating as a reference point

[20] Hazard rate is lower by -0.54-0.55+1.01=-0.08 for a win-stopper type after a win compared to baseline and lower by -0.55 after a loss. Therefore loss decreases the hazard rate much more than a win.

Covariate	Coef	exp(Coef)	p-value
Score	-0.62	0.54	< 0.005
Type	-0.65	0.52	< 0.005
Type × Score	1.13	3.08	< 0.005
Lag 1 score	0.04	1.04	< 0.005
Type × Lag 1 score	-0.04	0.96	< 0.005

Table 6: CPH with type heterogeneity and Lag 1 game

One possible factor to effect player's stopping decision could be the session's starting rating. If a player's stopping rule is to end a session once he reaches a higher rating compared to what he started with, then the rating difference between the first and last game of a session should be significant. To test this hypothesis, we included rating change since the start of the session in the CPH and estimated the model.

Covariate	Coef	exp(Coef)	p-value
Score	-0.61	0.54	< 0.005
Type	-0.67	0.51	< 0.005
Type × Score	1.10	3.01	< 0.005
Rating Change	0.00	1.00	< 0.005
Type × Rating Change	-0.00	1.00	< 0.005

Table 7: CPH with type heterogeneity and rating change since the start of the session

Table 7 shows that the change of the rating since the session started is not an important variable to explain stopping behavior, neither is interaction term between the type of the player and rating change. It is important to note that change in rating after every game is in the interval of [-12, 12], the coefficient for rating change is 0.002, and for rating change and interaction term is -0.001; therefore, the effect of rating change seem not essential for the stopping decision.

1.5 Conclusion

In this paper, we investigate stopping behavior in an environment free of monetary incentives and identify factors that determine how people make a stopping decision. We use rich data that provides information about an individual over the years in many occurrences. We identify two behavioral types: win-stoppers and loss-stoppers. Win-stoppers are more likely to stop playing after a win, while Loss-stoppers are more likely to stop playing after a loss. With conservative parameter values, we classify 84% of players as behavioral types, one-third of which are win-stoppers, and the rest loss-stoppers.

We develop a dynamic discrete choice model in which the agent may have time non-separable preferences. The model allows for the future game utility to depend on the outcomes of the current game and can capture the heterogeneity in stopping behavior. We use the model to estimate the model parameters and then test alternative market designs using counterfactual analysis. The results show that using type identification in the matching algorithm can increase session length by 5.44%.

Although the industry for online games is sizable, little is known about the determinants of length of play. In this paper, we document a robust behavioral force that drives players' stopping decision. The results in the paper might not be limited to a chess game, and the results can be applied to online games more broadly. Our approach can also be applied to other settings. For example, ride-sharing platforms such as Uber and Lyft can use a similar approach of looking heterogeneity to increase drivers' labor supply by using the information on the user's tipping behavior to match riders with drivers.

Chapter 2: Market Design for Gambling Platforms: The Effects of Biased Beliefs

Every time we receive new information, we update (at least, we are supposed to update) our beliefs. A crucial issue in dynamic environments is how people update their prior beliefs, i.e. how people process a piece of new information. In recent years psychologists and economists have challenged the Bayesian updating rule, the primary tool used by economists to model belief updating. Many studies show that people do not follow Bayes' rule; however, the significance and specific patterns of these deviations remain an open question.[21]

In this paper, we explore how people update their beliefs using investors' exit decisions. Proprietary data from online bookmaker shows that bettors' (investors') decision to cash-out their bet (investment) follows systematic patterns. In particular, we observe that bettor's choice to cash-out depends on (*i*) prior winning probability of a bet; (*ii*) whether new information arrived before cash out is good or bad; (*iii*) how good or how bad a piece of new information is. Unfortunately, we do not have access to the full data in the sense that we observe only cashed out bets and non-cashed-out bets are missing. To fill this gap, we design a new game, the "cash-out" game, and study it in the lab.

The cash-out game is a dynamic game that lasts up to two periods. In the first period, subjects decide how much to invest. In the case of positive investment, they move to the second period where they observe new information about their investment and have to decide to withdraw their investment or not. The investor chooses to withdraw when her belief about the success of her investment is not sufficiently high. However, if she invested in the first place, cashing out decision might be because of belief updating. To explore how people update their beliefs, we use findings from bookmaker data to have variation in the following three dimensions of the cash-out game. The first dimension is that updating might depend on a **prior** belief – updating of a more likely event might be different from updating a less likely event. The second dimension is that even if the priors are the same, **news direction** matters – one might over- or under-update compared to Bayesian prediction based on whether the news is good or bad. The third dimension is that, for a given prior and signal direction, updating might still be different for different **signal strengths** – we can rank both good and bad news by how good or how bad the news is.

In our experiment, the same subject faces the game with different priors; for each prior subject can receive good or bad news; we allow news to have different strengths. In other words, the subject first makes an investment decision, then new information arrives, which can be one of the following: worst, very bad, bad, good, very good, or excellent. After

[21] At the end of this section there is a short survey of contradictory findings.

receiving the news, the subject makes the second decision: whether to withdraw the investment or not. The amount that they can withdraw depends on their initial investment and new information. The better the news, the higher the cash-out offer. We offer subjects the expected value of the investment. By design the cash-out game allows us to explain in one single unified framework: the impact of priors, the effect of the news direction, and the signal strength on decision making. Further, unlike the prior literature, we use subjects' decisions to infer their beliefs instead of directly asking them what their beliefs are.

We use the Bayesian updating rule in combination with expected utility theory to derive baseline predictions for both risk-averse and risk-seeking preferences. Different risk attitudes produce different sets of predictions. We use those predictions to explain (rationalize) patterns that we see in our experimental data.

Our first step in analyzing belief updating is to check if dynamics or framing matter. In other words, do people make the same choices if they face the same decision problem as a simple one-period lottery instead of the second period of the cash-out game? To test that, we recorded all the trade-offs that subjects faced in the second period of the cash-out game and then had them face the same trade-off at the end of the experiment in a sequence of simple, static lottery questions. We find that 94% of subjects reversed their decision in at least one out of ten scenarios. Moreover, 43% of subjects made precisely the opposite decision in at least half of the cases, i.e., 43% of subjects favored a risky option (chose lottery over guaranteed outcome) in the static lottery question. In contrast, they preferred a safe option (decided to withdraw their investment) in the second period of the cash-out game or vice versa. We argue that such a substantial change is because people take into account the environment and do not make decisions only based on underlying lottery.

We document substantial heterogeneity on belief updating among subjects. Subjects use different decision-making strategies in the second period of the cash-out game. Two groups with the highest number of subjects have exactly the opposite strategy. The first group is more likely to withdraw their investment if they get better news, so have **"above"** strategy.[22] We say that subjects in that group under-interpret good news. The intuition is as follows. Investing in the first period means that expected utility from investing is at least as high as the expected value of investment (expected return from investment is its expected value). If a subject decides to cash out after good news, when again the expected value is offered, one of the explanations is that she did not update her initial belief as much as her actual probability of success increased, i.e. she under-interpreted good news. The second group is more likely to cash-out if they get worse news — they over-interpreted bad news. We say

[22] An agent has "above" strategy if she cashes out above some signal when signals are ordered according their informativeness about a winning state.

that they have **"below"** strategy.[23] Similar to the story of the first group, the intuition is simple. If a subject invests and decides to cash out after bad news, it implies that she updated her prior belief more than her actual chances decreased, i.e., she over-interpreted bad news. Moreover, we show that respectively 7.9% and 0% of subjects changed their cash out strategy from "below" to "above" and from "above" to "below" when they faced the same game with different initial probability. This shows that subjects are stable in their cash-out strategy when they face a new game, even when initial probability is different.

The belief updating literature has no consensus on this topic. Some papers show that individuals assign too much (Kahneman and Tversky (1973)) or too little weight (Grether (1992)) to the prior, others argue that changes in beliefs depend whether the news is good or bad (Eil and Rao (2011)), and some researchers claim that people over or undervalue information based on its quality (Ambuehl and Li (2018)). Our findings explain these mixed results.

Our first massage is that ignoring heterogeneity among subjects might result in contradictory results when one looks at average beliefs after the news arrival. If one experiment has more people who over-interpret bad news and an other has more people who under-interpret good news, depending on the portion of each group, results can substantially change. Second, we show that belief updating depends not only on news direction but also on how good or how bad a piece of news is. For example, Möbius et al. (2014) argues that subjects overweight good news, while Coutts (2019) found over-weighting of bad news. Those results are not necessarily contradictory if we take into account that good news in the former is "better" news than good news in later (similarly, the bad news is "worse" in former).

The main application of our results is in online gambling. We show that subjects are heterogeneous in their cash-out strategies. Bookmakers can use heterogeneity among bettors to design (improve) a cash-out scheme. First, they need to identify bettors type, weather a bettor is more likely to cash-out after better or worse news. This can be done by looking at bettor's history of cash-out strategies. We show that those strategies are stable over different initial probabilities, so bookmaker can use this information to predict bettor's future cash-out decisions. Second, bookmaker can extract more surplus from bettors by adjusting cash-out offer amounts according to bettor's type. For example, bookmaker can lower cash out offer after a worse news for bettors who are more likely to cash-out after a worse news. Similarly bookmaker can lower cash-out offer after a better news for bettors who are more likely to cash-out after a better news.

The rest of the paper is organized as follows. Proprietary empirical data is discussed in section 2. Baseline theoretical predictions are discussed in Section 3. Section 4 presents

[23] Subjects cash out "below" some signal.

the experimental design. Section 5 analyzes experimental data and shows the main results. Section 6 concludes.

2.1 Bookmaker Data

We use a bookmaker data to motivate our experimental design. The data is about online sports betting. Here is how online sports betting works. Bettor has access to odds of all possible events they can bet on. Odd is objective evaluation of probability of the event. Bettor constructs ticket by selecting the events that they want to bet on. Ticket can be an event or collection of events.[24] After bettor makes a selection, she has an option to "cash-out" (withdraw their bet) any time until the last event in their ticket is realized. If last event in the ticket is realized and ticket is not cashed out, bettor wins if all her predictions were correct and loses if even at least one of her prediction is incorrect.

Let's use example of a ticket with only single event to explain betting dynamics. Assume event is an outcome of a game between two soccer teams, Real Madrid C.F. and FC Barcelona. Let's say probability that Real Madrid wins is 0.50 and a bettor bet $5 for Real Madrid to win. Bookmaker always offers expected value in case of win.[25] Therefore bettor has 50% chance of winning $10 and 50% chance of losing her own stake.[26] After bet is made, anything that happens during the game is either good or bad news for bettor and affects to cash-out offer. For example if Real Madrid scores, objective probability that Real Madrid wins increases. Assume after scoring Real Madrid has 75% chance of winning, then bookmaker offers bettor $0.75*\$10+0.25*\$0 = \$7.50$ to cash-out. Contrary, if Barcelona scores and objective probability drops to 0.25 then bookmaker offers $0.25*\$10+0.75*\$0 = \$2.5$ bettor to cash-out. If the cash-out offer is accepted, bettor takes money and ticket is cancelled.

In our data one observation is one cashed-out ticket with the following information: stake (how much was bet), initial probability of winning, probability of winning at the time of cash out, and cashed out amount. All the observations are bets that has been cashed out. Unfortunately, we do not have bets that were not cashed out. We divide observations into two groups, cash out after good and after bad news. Division rule is straightforward. If person cashed out more (less) then her stake, something good (bad) must have happened for her tickets to gain (lose) value.[27] In total we have 4520 observations, with 2133 observations

[24]The ticket can be a collection of facts about different games. It is not mandatory to have one game in each ticket, for example bet can be made about which team starts the game, or how many goals will be scored in total or even who will eat pie during the game (this odd was offered for one particular game http://www.dailymail.co.uk/sport/football/article-4243656/Sutton-reserve-goalkeeper-Wayne-Shaw-eats-pie.html)

[25]Bookmaker keeps 3-7% for itself.

[26]Bettor buys lottery with expected value of 0.5*$10+0.5*$0 for $5.

[27]Alternatively, we can compare initial and updated probability to classify observations into cash out

cashed out after a bad news and 2387 observation after a good news.

Initial and updated probability of the ticket provides a nice dynamic environment where we observe bettors decision twice for the same ticket, before and after the news. To quantify signal strength (informativeness) that links initial and updated probabilities, we use the Bayesian updating rule with the following signal structure. Assume, there are two possible states of the world $\omega \in \Omega = \{W, L\}$ with initial probability of state W is p_0. Before actual state of the world is observed, binary signal (new information) $s \in S = \{g, b\}$ arrives about the state of the world. Assume signal is symmetric, $Pr(g|W) = Pr(b|L) = \alpha \geq 0.5$. We can use the Bayes rule to find updated probability after each signal:

$$p_{1g} \equiv Pr(W|g) = \frac{\alpha * p_0}{\alpha * p_0 + (1-\alpha) * (1-p_0)} \tag{10}$$

$$p_{1b} \equiv Pr(W|b) = \frac{(1-\alpha) * p_0}{(1-\alpha) * p_0 + \alpha * (1-p_0)} \tag{11}$$

Since we observe initial (p_0) and updated probabilities of winning after good (p_{1g}) and bad (p_{1b}) signals, we use equations 10 and 11 to find signal strength (α) for each observation in our data. Equipped with news quality (signal strength) measure, we analyze effects of new information on cash-out decision. Signal strength α calculated from equations 10 and 11 is continues variable from 0.5 to 1. Closer to 0.5 less informative signal is, i.e. initial and updated probabilities are close to each other. For data visualisation let's divide $[0.5, 1]$ into grids of size 0.05. Figure 8 presents cash out frequencies for each grid of signal strength.

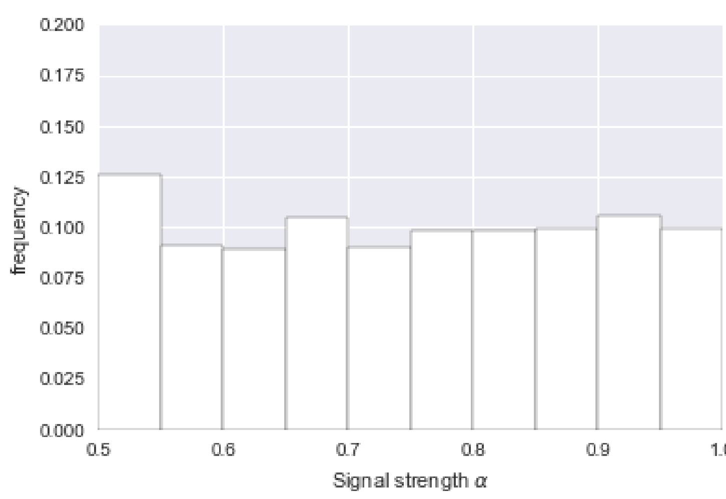

Figure 8: Relationship between signal and cash-out frequency

after good and bad news. If winning probability increased (decreased), it is because something good (bad) happened that increased (decreased) winning chances.

Figure 8 shows that cashing out is more or less uniform across different signals. In other words, cash out decision does not depend on how informative signal is. However, if we separate good and bad news cases we see a huge difference.

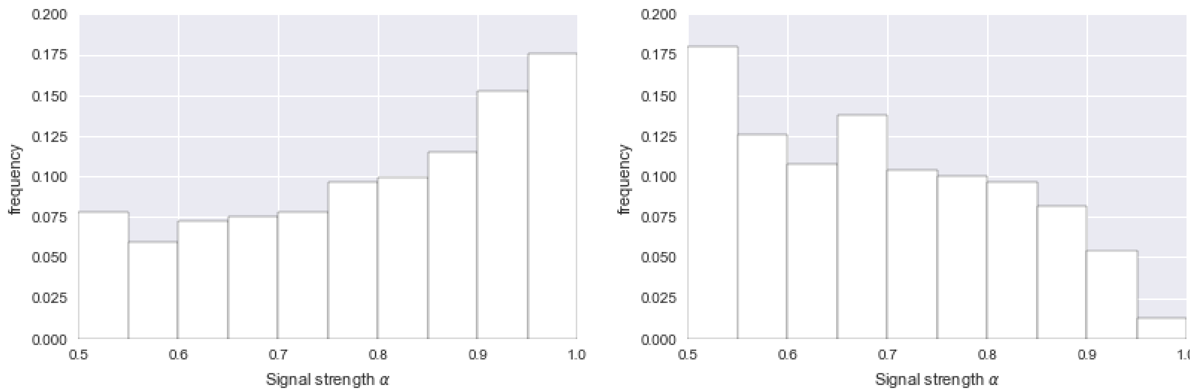

(a) Good News

(b) Bad news

Figure 9: Relationship between signal and cash-out frequency for good and bad signals

Figure 9 shows relationship between signal informativeness and cash out frequencies separately for good and bad signals. Here we see that there is asymmetry which need to be explored deeper. Figure 9a shows that bettors are more likely to cash-out after a good news when α is higher, i.e. when they receive more informative good news. The opposite is true for bad news, bettors are more likely to cash out if bad news is less informative (figure 9b).

Another asymmetry that we see in cash-out decisions is in the relationship between initial probability of winning and cashed out signal strength for good and bad signals (figure 10). Figure 10a shows that median and average cashed out signal strength decreases as initial probability of winning increases for good news side, and figure 10b shows that the opposite is true for bed news side.

As mentioned earlier we have data only for cashed out tickets and bets that are not cashed out are missing. To fill this gap, we use three main findings from observational data to design a laboratory experiment. First, cash-out decision depends on initial probability. Second, there is substantial asymmetry in cashing out after good and bad news. Third, cash out frequency depends on signal strength, i.e. how good or how bad a signal is affects cash-out decision. Experiment presented in section 3.5 has variation in all three dimensions.

2.2 Theoretical Framework

In the following section, we use the Bayesian updating rule in combination with expected utility theory to derive baseline predictions for both risk-averse and risk-seeking preferences.

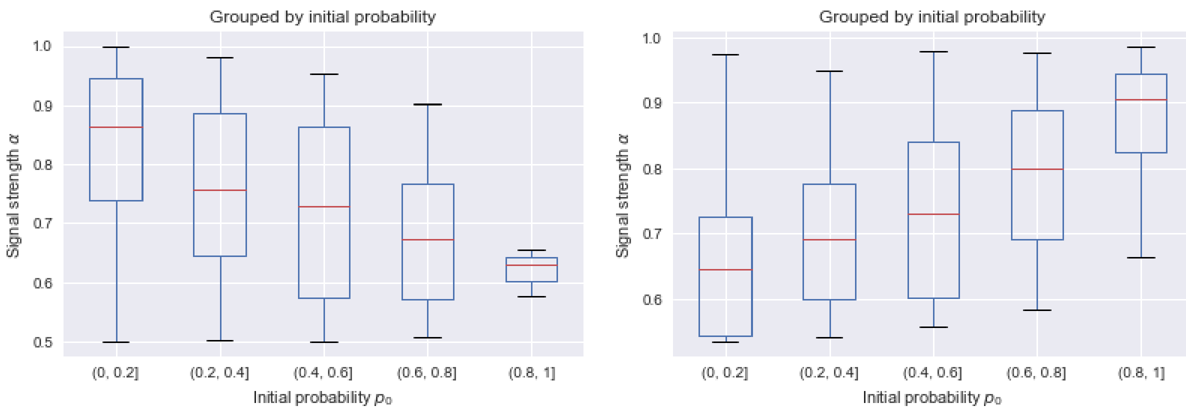

(a) Good News (b) Bad news

Figure 10: Relationship between initial probability and signal

First, we define objective probabilities and agents' beliefs in section 2.2.1, then we move to section 2.2.2 to discuss the decision-making process of an expected utility-maximizing agent in the two-period investment "Cash-Out" (CO) game.

2.2.1 The Model Primitives

Consider an agent who forms and updates beliefs over two possible states of the world $\omega \in \{\mathbf{W}in, \mathbf{L}ose\}$. The nature selects $\omega = W$ with prior probability $p_0 = Pr(\omega = W)$. An agent receives signal $s \in \{s_1, s_2, ..., s_n\} \equiv \mathcal{S}$ and $|\mathcal{S}| = n$. Assume signals are ordered by how much it increases probability of winning state, i.e.

$$s_1 \prec s_2 \prec ... \prec s_n \iff Pr(W|s_1) < Pr(W|s_2) < ... < Pr(W|s_n)$$

We let $p_{1i} = Pr(\omega = W|p_0, s_i)$ to be the Bayesian posterior probability of winning after observing signal s_i. Define α_i^W and α_i^L to be probability of observing signal s_i respectively when the correct state is W and L i.e. $\alpha_i^W = Pr(s_i|W)$ and $\alpha_i^L = Pr(s_i|L)$.[28] Finally, let's denote p_{0i} to be probability of observing signal s_i.[29] Now we have all we need to present the Bayes rule for updating objective probabilities:

$$p_{1i} = \frac{\alpha_i^W * p_0}{\alpha_i^W * p_0 + \alpha_i^L * (1 - p_0)}$$

All the probabilities that we defined so far are objective probabilities. Agents' beliefs are not restricted to the same objective probabilities. Thus, define μ_0 and μ_{1i} to be an agent's

[28]Two obvious conditions that α's satisfy are $\sum_{i=0}^{n} \alpha_i^W = 1$ and $\sum_{i=0}^{n} \alpha_i^L = 1$.
[29]Probability of observing signal s_i is $p_{0i} \equiv Pr(s_i) = \alpha_i^W * p_0 + \alpha_i^L * (1 - p_0)$.

initial and updated (after observing signal s_i) beliefs that the realised state is W. Also, we allow agents to have different beliefs on the probability of observing each signal in each state and denote by $\beta_i^W = Pr(s_i|W)$ and $\beta_i^L = Pr(s_i|L)$.[30] We use denote an agent's belief of observing signal s_i by μ_{0i}.[31] Similar to objective probabilities, an agent's prior and posterior beliefs can be connected using the Bayes rule:

$$\mu_{1i} = \frac{\beta_i^W * \mu_0}{\beta_i^W * \mu_0 + \beta_i^L * (1 - \mu_0)}$$

Since our goal is to infer an agent's beliefs form her decisions we need to have a notion of utility. Define an agent's utility function to be continues and increasing $u(\cdot)$ and we assume agents maximize their expected utility.

2.2.2 Two Period Discussion

Subject's play a two-period single-player investment game – cash-out game. In the first period, an agent faces investment opportunity, which has an initial probability of p_0 to be successful. If she invests any positive amount $x > 0$, she gets $\frac{x}{p_0}$ with probability p_0 and 0 with probability $1 - p_0$ at the end of the game (second period). However, before the outcome of an investment is realized, the agent observes a public signal s_i (new information) about her investment. New information changes the probability of success from p_0 to p_{1i} and an agent is offered to withdraw (cash-out) her investment. Cash-out offer amount depends on signal, better the signal, higher the cash-out offer, to be more specific, $\frac{p_{1i} * x}{p_0}$ which is exactly the expected value of an investment.[32]

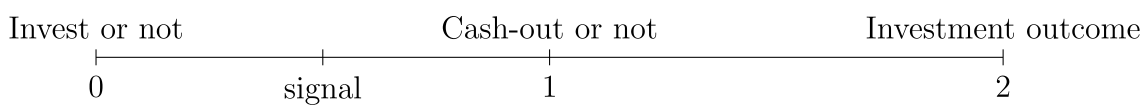

Figure 11: Timeline

We have all the notations that we need to express trade-off that an agent faces at the beginning of the game – invest or not. On one hand, if she does not invest she gets utility from x. However, if she does invest she receives signal $s_i \in \mathcal{S}$ and will make another decision, to cash-out or not.[33] Therefore, an agent invests if the following condition holds.

[30] Similar to α's, we have similar conditions for β's $\sum_{i=0}^{n} \beta_i^W = 1$ and $\sum_{i=0}^{n} \beta_i^L = 1$.

[31] An agent's belief of observing signal s_i is $p_{0i} \equiv Pr(s_i) = \alpha_i^W * p_0 + \alpha_i^L * (1 - p_0)$.

[32] This is the amount that an agent needs to invest if she joins at the end of the first period to get p_{1i} chance of receiving $\frac{x}{p_0}$. Algebra: $\frac{p_{1i} * x / p_0}{p_{1i}} = \frac{x}{p_0}$.

[33] We assume in case of equality an agent choose a risky option both in the first as well as in the second period, this assumption does not affect our results.

$$u(x) \leq \sum_{i=0}^{n} \mu_{0i} * \max\Big\{u\Big(\frac{x}{p_0} * p_{1i}\Big),\ \mu_{1i} * u\Big(\frac{x}{p_0}\Big)\Big\} \tag{12}$$

One the left of equation 12 we have utility from not investing. On the right-hand side, there is the expected utility from investing. As discussed earlier, if an agent invests, she receives signal $s_i \in \mathcal{S}$ with probability μ_{0i}. After observing the signal she makes a second decision, to choose maximum between cash-out utility (the first term inside max operator) and expected utility from keeping the investment.

To make discussion easier to follow, let's assume we have two sets, a set of signals which leads to **C**ash-**O**ut denoted by $\mathcal{CO}$ and a set of signals which leads to **N**ot **C**ash-**O**ut denoted by $\mathcal{NCO}$. Since each signal can only lead one and only one decision, we have $\mathcal{CO} \cup \mathcal{NCO} = \mathcal{S}$ and $\mathcal{CO} \cap \mathcal{NCO} = \emptyset$. Definition implies:

$$s_i \in \mathcal{CO} \Longleftrightarrow u\Big(\frac{x}{p_0} * p_{1i}\Big) > \mu_{1i} * u\Big(\frac{x}{p_0}\Big)$$

$$s_i \in \mathcal{NCO} \Longleftrightarrow u\Big(\frac{x}{p_0} * p_{1i}\Big) \leq \mu_{1i} * u\Big(\frac{x}{p_0}\Big)$$

We can re-write equation (12) as,

$$u(x) \leq \sum_{i \in \mathcal{CO}} \mu_{0i} * u\Big(\frac{x}{p_0} * p_{1i}\Big) + \sum_{j \in \mathcal{NCO}} \mu_{0j} * \mu_{1j} * u\Big(\frac{x}{p_0}\Big) \tag{13}$$

Now we can make assumptions on signals and utility functions to draw some predictions. We move all the derivations in the Appendix B Derivations. Let's start from the most obvious one:

Claim 1-a: *assume $p_0 = \mu_0$, $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, and $u(\cdot)$ is* ***concave****, then investor never invests and always cashes out.*

Claim 1-b: *assume $p_0 = \mu_0$, $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, and $u(\cdot)$ is* ***convex****, then investor always invests and never cashes out.*

The first two claims say that if an agent follows objective probabilities, then her decision completely depends on the shape of her utility function. If she is risk-averse (concave utility), since we offer expected value for investment return, she takes a guaranteed amount in the first period (so not invests). If she invests for some reason, she cashes out after the first period for every signal. Similarly, if an agent is risk-seeking (convex utility function), she always takes a risky option. Invests and does not accept cash-out offers.

Now we relax condition $\mu_0 = p_0$ and allow agents to have arbitrary initial beliefs. We still assume that an agent is Bayesian and her beliefs on possibility of each signal equals to objective probabilities i.e. $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$.

Claim 2-a: *assume $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, and $u(\cdot)$ is* ***concave****, if an agent is Bayesian and cashes out after some signal $s_i \in \mathcal{S}$, then she cashes out after any signal $s \in \mathcal{S}$, such that $s \prec s_i$. Also, if an agents does not cash out after $s \in \mathcal{S}$, she does not cash out after any $s \succ s_i$.*

Claim 2-b: *assume $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, and $u(\cdot)$ is* ***convex****, if an agent is Bayesian and cashes out after some signal $s_i \in \mathcal{S}$, then she cashes out after any signal $s \in \mathcal{S}$, such that $s \succ s_i$. Also, if an agents does not cash out after $s \in \mathcal{S}$, she does not cash out after any $s \prec s_i$.*

Claim 2 suggests that a Bayesian agent who does not over- or under- interpret any signals has a threshold strategy. If she has a concave (convex) utility function, she cashes out worse (better) signals more likely. Claim 2 gives us baseline predictions for Bayesian agents with both risk-averse and risk-seeking preferences.

2.3 Experimental Design

The experiment was designed to follow the spirit of the empirical data in order to study the effects of new information on decision making. We use both within-subject and between-subjects comparisons to analyze changes in beliefs in the face of new information.

All experiments were conducted at the Laboratory for Economics, Management and Auctions at Penn State University using oTree (Chen et al. (2016)). Subjects were recruited from the general undergraduate population, and each subject participated in only one session. We conducted 6 experimental sessions in the spring of 2019. Treatment 1 and treatment 2 each had 3 sessions, with number of subjects respectively 21, 25, 25 in treatment 1 and 13, 23, 24 in treatment 2. In total there were 131 subjects. The average (median) payment from the experiment was \$9.68 (\$10), plus a \$7 participation fee. On average, sessions lasted about 1 hour.

The study consisted of four parts. The first two parts were the “cash-out” games, in the third part we tested subjects’ abilities using raven matrices. In the last part, we asked 18 lottery questions that corresponded to the lotteries subjects faced in the second stage of the two cash-out games. We had two treatments, where the only difference was the order of the first two parts. The main part of the experiment was the cash-out game

2.3.1 The Cash-Out Game

We develop the cash-out game specifically for this experiment. Subjects are told that they have an opportunity to invest in a lemonade stand. The lemonade stand can operate for up to 10 days, in groups of 5 days at a time. For the lemonade stand to be successful, in the first cash-out game, at least 5 of the days it operates have to be sunny. For the second cash-out game, at least 6 of the days it operates have to be sunny. Subjects know the minimum number of sunny days they need for success at the start of a Round before they make any decisions. As an example, we consider the case where subjects need to operate the lemonade stand at least 5 sunny days for their investment to be successful.

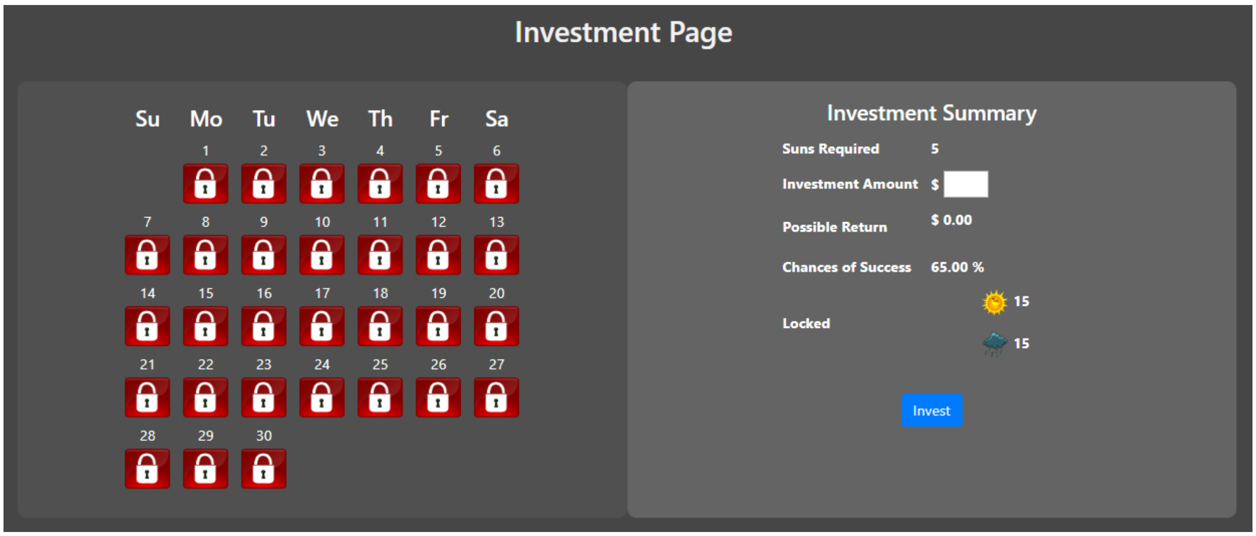

Figure 12: Investment stage.

Figure 12 represents a screenshot of the first screen that subjects face. The left-hand side of the image displays a 30-day calendar. Out of those 30 days, 15 hide suns and 15 hide clouds. Subjects are given a \$15 endowment and have to decide how much of their \$ 15 endowments they wish to invest in the lemonade stand. They can invest the following amounts: \$0 (i.e. not invest), \$1, \$5, \$10, and \$15. Subjects input one of those numbers in the "Investment Amount" field, which automatically changes the next field "Possible Return". As mentioned earlier, we always give subjects the expected value of an investment, so a possible return is calculated as the ratio of the invested amount with the probability of success. Information about returns is available to subjects, and they can also input different numbers to see different possible return amounts.

If a subject invests zero, the game is over and she keeps her entire \$15 endowment. If the subject invests a positive amount, her next decision is to unlock 5 days from the calendar by clicking lock symbols. Figure 13 shows the screen that a subject faces after investing \$10 (shown under "Investment Summary"). Green unlocked symbols represent the 5 days that the subject selected to operate the lemonade stand.

After selecting five days, a subject observes the weather forecast for those days and based

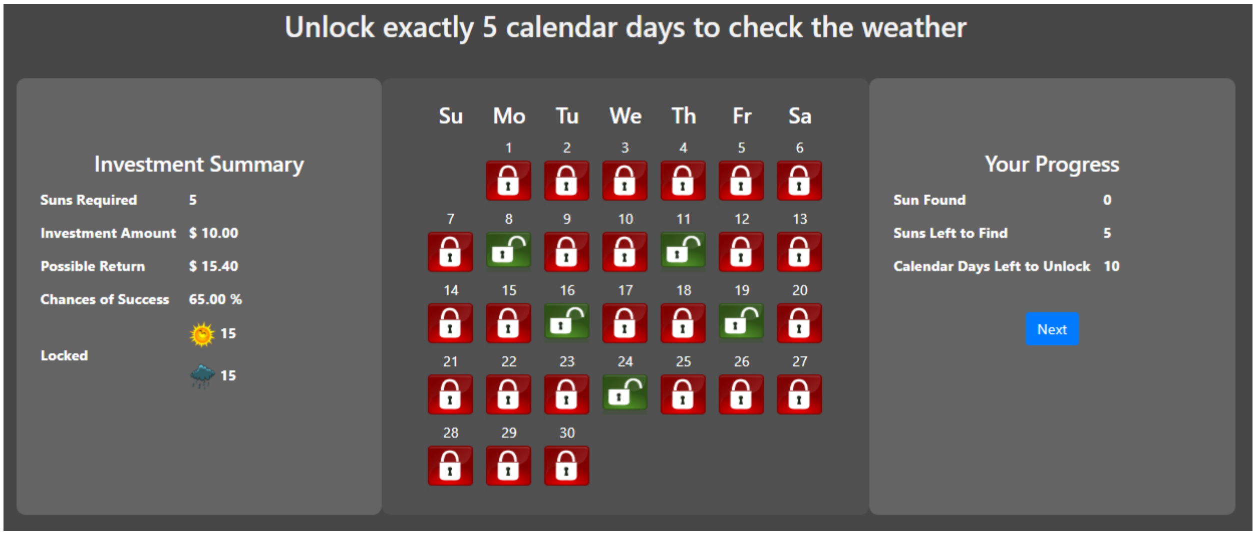

Figure 13: Subject selects 5 days to operate a lemonade stand.

on the number of suns, a cash-out offer is made. Figure 14 presents the case when 3 out of selected 5 days are predicted to be sunny. Since success probability is increased compared to the start of the game, the subject is offered more than she invested: $12.55 to withdraw her investment.

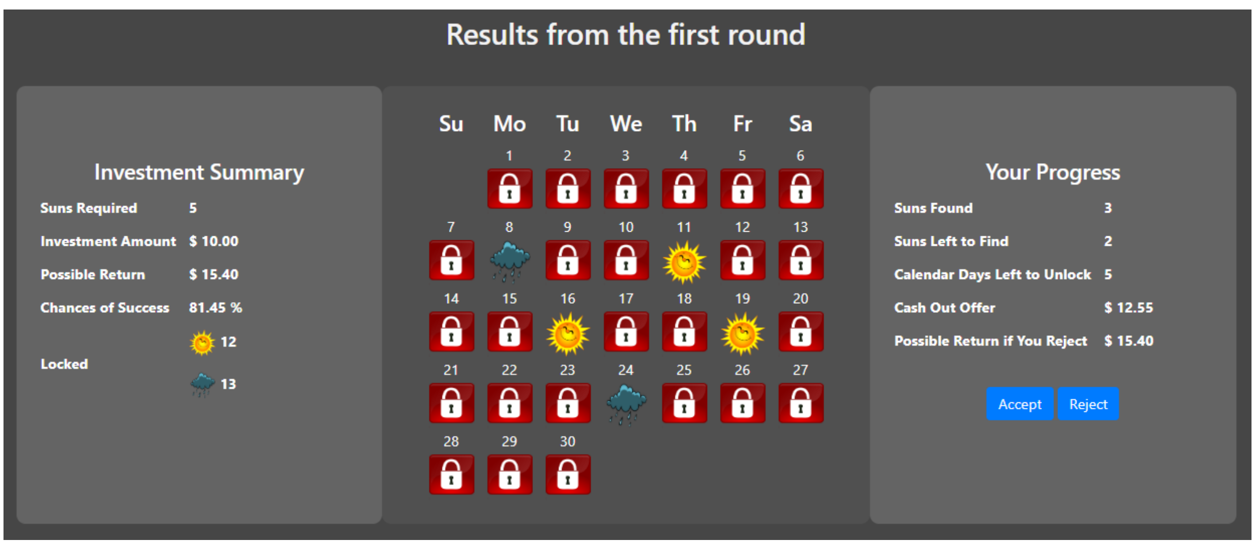

Figure 14: Subject receives cash-out offer based on revealed numbers of suns.

If a subject accepts the cash-out offer, the game is over. In case of rejection, a subject selects 5 days from the remaining 25 days and then observes the final outcome. If the number of sunny days is at least 5 out of 10 selected days, the subject receives "Possible Return" ($15.40 in this case) plus whatever she saved by not investing ($5 in this case). If the number of sunny days is less then 5, then the invested amount ($10 in this case) is lost and the subject receives only the saved amount.

2.3.2 Experimental Procedure

In treatment 1, the first part of the experiment was the cash-out game, where the required number of suns for the investment to be successful was 5. After instructions were read, subjects played 15 rounds of cash-out game for practice to learn the game and try different strategies. It was followed by a final round. In the final round, subjects again made an investment decision, and selected 5 days on the calendar if they invested positive amount. We then used the strategy method to elicit their decisions for all possible cases. We used the strategy method to avoid asymmetric data.[34] Subjects were informed about the strategy method in advance and they faced all possible numbers of revealed suns in random order. Here is how we described the strategy method to the subjects:

> *Out of the 5 days you selected, none may be sunny, one may be sunny, two may be sunny, three may be sunny, four maybe sunny or all may be sunny. We will not tell you (yet) which of these actually occurred. Instead, at this point you will be asked to tell us what you would do in each of the following possible Scenarios: if you uncovered 0 suns, 1 sun, 2 suns, 3 suns or 4 suns. We will not show you a scenario of the case where you uncovered 5 sunny days in Stage 1 since in that case your lemonade stand is already successful and you will earn the returns associated with success. Your choice for each scenario is whether or not to accept a "cash-out" offer. After you have made your decision for each of the scenarios, we will tell you which one actually occurred, and it is your decision in that scenario that will be implemented.*

In the second part of the experiment subjects also played the cash-out game, but in this case, they had to find 6 suns for their investment to be successful.

We used raven matrices in the third part to divert subjects' attention from the decision they made in the first two parts and to asses their abilities. In the last part of the experiment, we wanted to evaluate whether decisions in the cash-out game were consistent with decisions when subjects faced the same trade-off but without the dynamic aspect of the game (i.e. updating information). To do that, we recorder all the trade-off that subjects faced during the first two parts of the experiment and presented them as simple lottery questions .[35] [36]

[34] Without the strategy method, approximately 68% of our observations would have been for either 2 or 3 found suns in the first stage.

[35] For example, the corresponding lottery question of figure 14 would be trade-off between 100% chance of receiving \$12.55 versus 81.45% chance of receiving 15.40 and 19.55% chance of receiving \$0.

[36] Anyone who invested positive amount in both parts, faced 10 cash-out scenarios. Also, we asked them to decide between guaranteed \$1, \$5, \$10, or \$15 and possible returns with their respective probabilities.

At the end of the experiment subjects filled out an optional questionnaire. For treatment 2 we reversed the order of the first two parts. Subjects in treatment 2 played the cash-out game with 6 required suns in the first part and with required 5 suns in the second part. The rest of the experiment was the same for those two treatments.

2.4 Data Analysis and Results

We start by presenting aggregate investment data. Figure 15 shows investment frequencies for 5 and 6 sun cash-out games. Around 22% of subjects did not invest anything each case. Since the 5 suns case is less risky (probability of success 0.65) then the 6 suns case (probability of success 0.35) we expect less investment for the riskier case. For the 5 suns case the most popular investment is either $10, while for 6 suns case it is $5. [37]

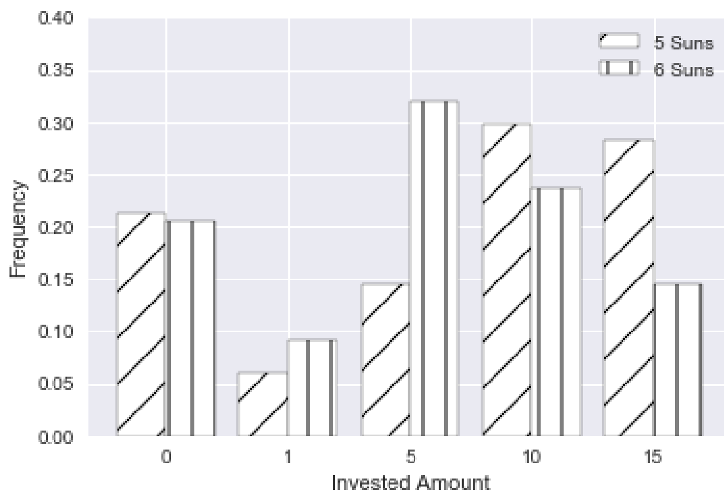

Figure 15: Investment frequencies.

2.4.1 Results

The first result (section 2.4.1) argues that there is history (case) dependence in decision making. In other words, we show that subjects make different choices if they face the exactly same decision problem as a simple one-period lottery instead of the second period of the cash-out game. The second result (section 2.4.1) demonstrates the substantial heterogeneity among subjects. We divide subjects into multiple mutually exclusive groups based on cash

For example, guaranteed $10 VS 35% chance of winning $28.60. If a subject did not invest in any part of the study we asked lottery questions as if she invested $5 and faced respective cash-out trade-offs.

[37] Figure 37 in the appendix Appendix B Graphs by treatment shows that investment amounts are very similar across the two treatments, indicating that whether a subject see the 5 or 6 sun games first had little impact.

out strategy that they use. Interestingly, the two groups with the highest number of subjects have exactly the opposite strategy.

History Dependence

Do the dynamics or formulation matter? Put differently, do people look only at numbers when they make decision? The answer is yes for the first and no for the second question. Let's propose the following hypothesis:

H_0: History/Framing has no effect on decision making.

H_0 implies that agents' decisions are independent of a history or an environment as far as the underlying lottery is the same. If it is true then it does not matter whether we ask the same lottery question as second period cash-out game scenario or simple, one period lottery question. To test H_0, we recorder all the trade-offs subjects faced during the cash-out game and asked them separately as simple, one period lottery questions.

Difference between choices made in cash-out and lottery questions is substantial. We find that 94% of subjects reversed their decision in at least one out of ten scenarios. Moreover, 43% of subjects made exactly the opposite decision in at least half of the cases, i.e. 43% of subjects favored a risky option (chose lottery over guaranteed outcome) in the static lottery question, while they favored safe option (chose to withdraw their investment) in the second period of the cash-out game or the other way around.

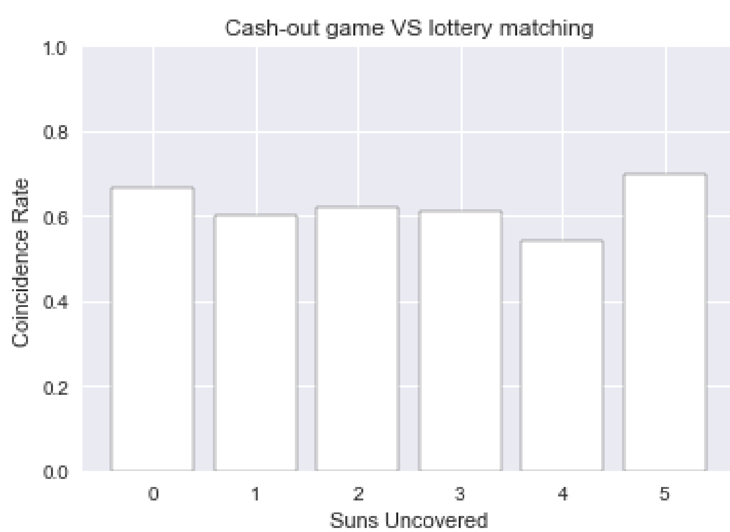

Figure 16: Matching the answers for the cash-out games and respective lottery questions.

Let's take a look if there is difference between coincidence of the cash-out game and the lottery questions across the signals. Figure 16 shows that for all signals, around 60% of the decisions coincide. Even if we separate cases for 5 and 6 suns it is hard to make any conclusions (figure 38 in the appendix Appendix B Graphs by treatment). However, there is one dimension that can give us some intuition for why subjects reverse there choice when they face lottery instead of the cash-out game.

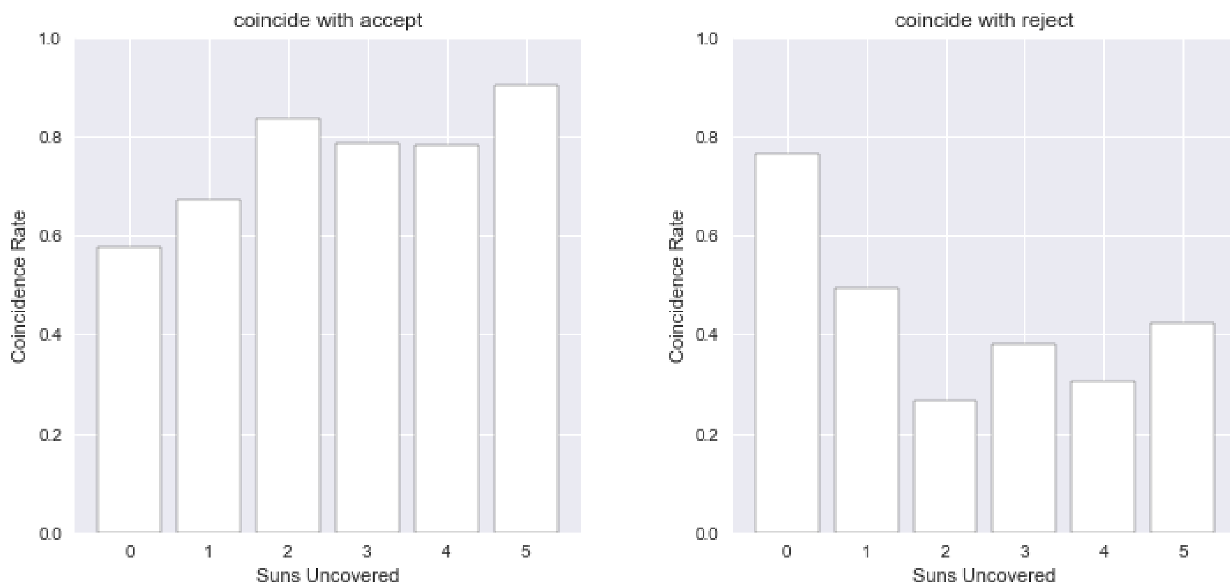

Figure 17: Matching the answers for the cash-out games and respective lottery questions separately for accepted and rejected cash-out scenarios.

Figure 17 shows that reversal happens mostly for rejected offers. In other words, if a subject accepts a cash-out offer, she is very likely to select guaranteed amount in the lottery, i.e. select safe option in both, the cash-out game as well as simple, one period lottery. However, for a rejected cash-out offers we have reversal. Subject is more likely to select safe option in the lottery after selecting a risky one in the cash-out game. We offer two explanations of why we have more reversal for rejected compared to accepted cash-out offers: framing and history.

Framing:

We argue that subjects might have an illusion that they can control the outcome. The illusion of control defined as the tendency for people to overestimate their ability to control events is documented both in lab (Langer (1975)) as well as in field data (Qadri and Shabbir (2014)). In addition, Elfenbein et al. (2017) shows that illusion of control is positively correlated with asymmetric updating which favors a good news. In our setting, if subjects have the illusion of control we should see it in the cash-out game since they are the one's who select the days in the game. In the lottery questions we do not expect any illusion of control because they have no intervention in outcome realization. To sum up, is a subject selects a safe option in the second period of the cash-out game, we expect her to select safe option in the lottery question as well (increase in belief because of illusion is not high enough). However, if a subject selects a risky option in the cash-out game it might be because of illusion of control, which she does not have in simple lottery questions.

History:

Empirical evidence suggests that risk attitudes often depend on an individual's history of gains and losses. Paper by Imas (2016) unifies mixed finding from the literature and claims that if a loss is not realized, individuals take on greater risk. For the worst signal we see less risk taking, but for all other signals we see that subjects take on higher risk compared to simple lottery question.

Heterogeneity

We start this section by looking at cash-out rates after each signal separately for 5 and 6 suns case. Figure 18 shows percentage of offers that has been cashed out for each signal. Without going any deeper we can argue that across signals average cash out behaviour is not very different. However the key word here is average, without taking into account heterogeneity among subjects. For example if we have two subjects one who cashes-out only bad and other cashes-out only good news, we will see that on average cash-out behaviour does not depend on signal strength, which would obviously inaccurate.

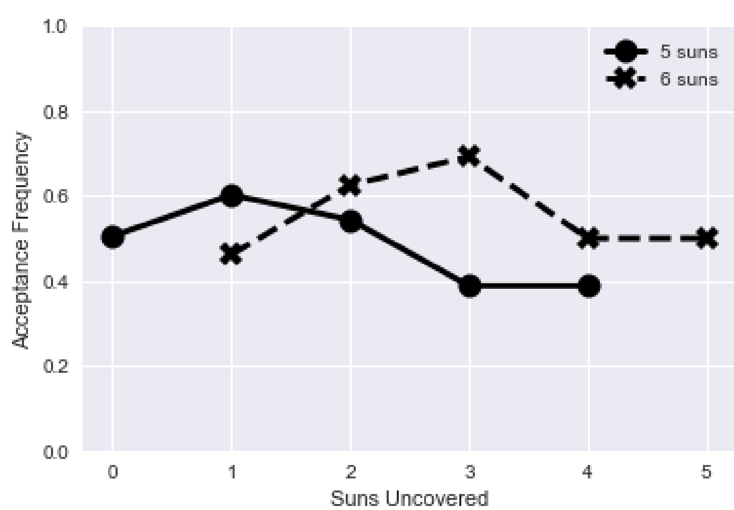

Figure 18: Cash-out frequencies after each signal.

Let's dive deeper to see if the example explained above have any support in the data. We look at each subject separately to investigate if they follow any specific cash-out rule. Figure 19 excludes all non-investors and describes portion of people who use respective cash-out rule. Let's describe all the rules. If subject with the "Above" rule, cashes-out after some signal $s \in \mathcal{S}$, she cashes out after any signal $s_j \in \mathcal{S}$ if $s_j \succ s$. The "Below" rule is reverse, if a subject cashes-out after some signal $s \in \mathcal{S}$, she cashes out after any signal $s_j \in \mathcal{S}$ if $s_j \prec s$. In other words, subjects with "above" rule cash-out above some signal and with "below" rule cash out below some signal. Above and below are used according of signal ordering from the worst to the best. Intuitively, rules "Never" and "Always" represents subjects who respectively never of always cash out. Subjects who follow the "Hump" rule cash out after less informative with higher chance. We put subjects who do not fit in any of the above mentioned category in "No rule" category.

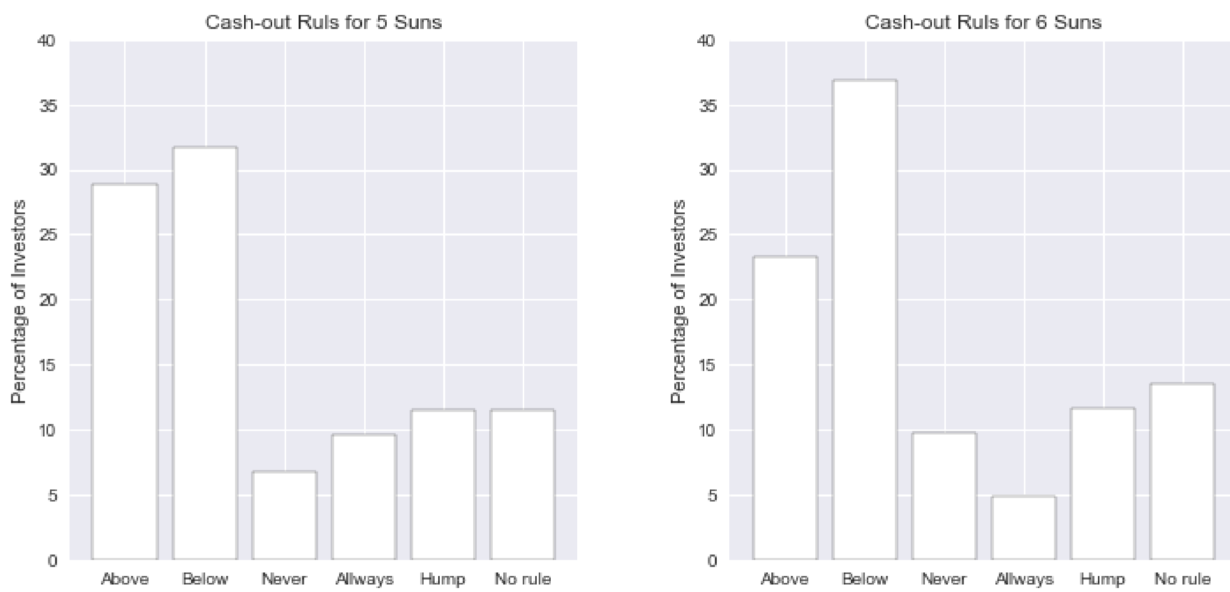

Figure 19: Percentage of investors who use respective cash-out rules.

Two rules that stands out the most are "Above" and "Below". Figure 20 presents cash out frequencies for both rules. It is obvious that there is substantial heterogeneity. Two largest groups have exactly the opposite cash-out strategy. The first group is more likely to withdraw their investment if they get better news. We say that subjects in that group under-interpret good news. Intuition is straightforward. Investing in the first period means that expected utility from investing is at least as high as the expected value of investment (expected return from investment is it's expected value). If a subject decides to cash out after good news, when again expected value is offered, one of the explanations is that she did not update her initial belief as much as her actual probability of success increased, i.e. under-interpret good news. The second group is more likely to cash-out if they get worse news – over-interpret bad news. Like above, intuition is simple. If a subject invests and decides to cash out after bad news, it implies that she updated her prior belief more than her actual chances decreased, i.e. over-interpret bad news.

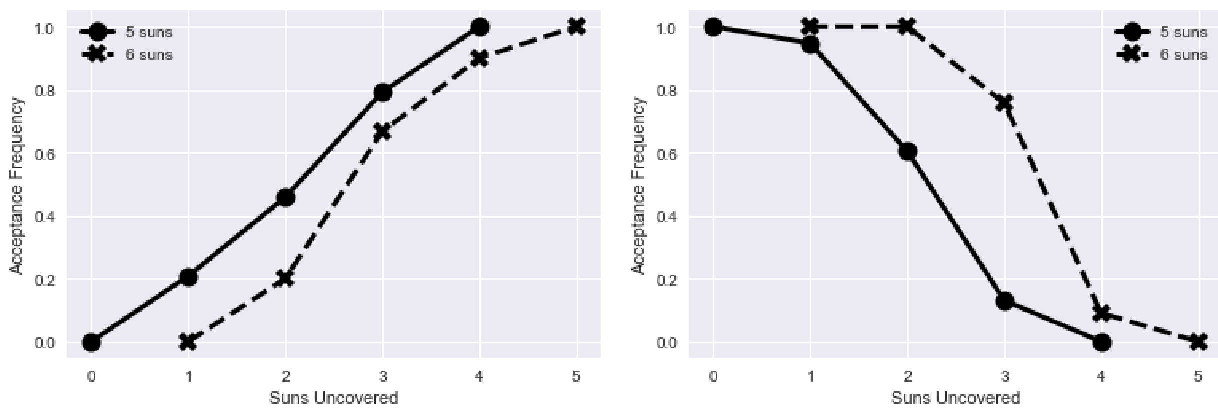

(a) The "Above" rule

(b) The "Below" rule

Figure 20: Acceptance frequencies after each signal for the "Above" and "Below" rules.

We use figure 21 to show that subjects who use a given cash-out rule for 5 suns are most likely to stick with the same rule for 6 suns case, irrespective of whether they played 5 suns case first or second. Two rules followed by the most subjects seems to be the most stable. Also subjects are very stable in their decision not to invest ("nobet" rule on the figure 21).

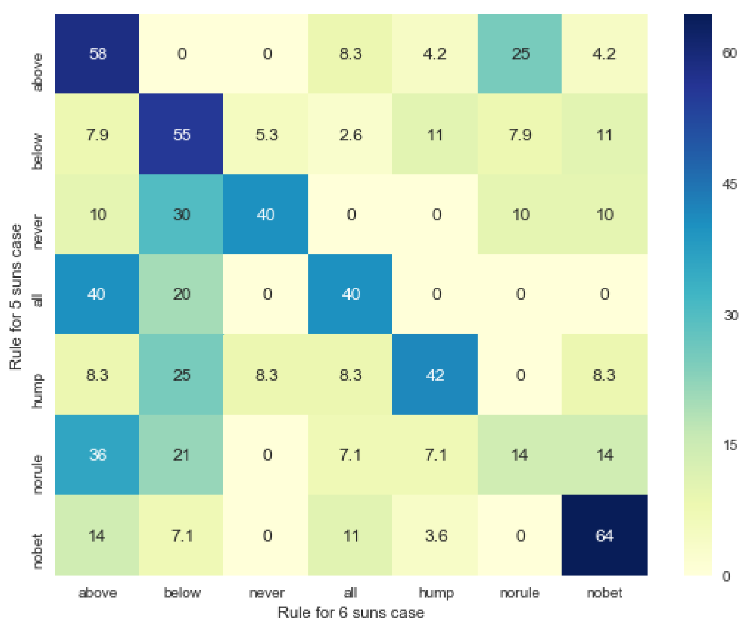

Figure 21: Percentage of investors who use respective cash-out rules.

As we discussed at the beginning of this subsection, if we have two big groups who follow exactly the opposite strategy, in the data it will appear as signal quality does not have substantial effect on decision (figure 18). We showed on figure 20 that is exactly the case and we have two large groups who use opposite strategies and signal informativeness has a huge effect on subjects' decision. Moreover, the same signal can have positive effect on one subjects cash-out decision while having negative impact on another subject.

2.5 Conclusion

In this paper, we investigate how people respond to the arrival of new information. We design experiment by taking into account findings from proprietary bookmaker data. In particular, we check how subjects react to new information for different prior beliefs, signal directions, and signal qualities. The main part of experiment is the cash-out game which is dynamic game, played by subjects for different initial beliefs, where they can receive good as well as bad signals and each good or bad signal can be classified as worst, very bad, bad, good, very good, or excellent. To check if dynamics (history) or framing has any effect on decision making we recorder all trade-offs subjects faced in the cash-out game and gave the same trade-off as a one period, simple lottery.

We have two main findings. First history/framing matters: 43% of subjects reversed their decision in at least half of the cases, when they faced the same trade-off in a simple one period lottery instead of second period of cash out game. We argue that observed phenomenon can be explained either by the "illusion of control" or by dependence of a subject's risk attitude on history of gains and losses. Second, we find heterogeneity in response to the arrival of new information. If we look at only average cash out frequencies, one might conclude that since probability of cash out is the same after each signal, signal informativeness has no effect on cash out decision. We show that such conclusion is inaccurate and signal informativeness is the most important determinant for cash-out decision. The reason we do not see it in average frequencies is that we have two main groups who have exactly the opposite cash-out strategies, one group cashes out after better signals and other cashes out after worse signals. Ignoring such heterogeneity can result in mixed findings that we see in the literature on the effects of good VS bad news.

Direct application of our findings is in online betting industry. After making their initial bet, bettors are allowed to cash-out their bet after any news with respective amounts. Although the industry for online betting is sizable, our findings can be applied to financial markets as well, where investors decide to buy and sell their financial assets after receiving new information.

Chapter 3: Collaboration Cycles of R&D Teams: Theory and Experiment

In this paper I investigate the possibility of collaboration among two firms, who consider whether or not to create a joint venture (JV).[38] A JV is a business arrangement in which two or more parties agree to pool their resources for the purpose of accomplishing a specific task.[39] Why is this important? While empirically there are many JVs, managing them is proved to be challenging. It is shown that 40% to 70% of JVs fail.[40]

In this paper we concentrate on analyzing Joint Ventures on R&D, so called Research Joint Ventures (RJV). It is encouraged in a lot of countries to create RJV's. For example, EU had two big programs EUREKA and Program Frameworks for Science and Technology (PFST) to sponsor RJVs. Also, in USA National Cooperative Research Act reduces the potential antitrust liabilities for research joint ventures to motivate companies to create joint research units.

To understand possible dynamics of RJV's, I use a dynamic model with asymmetric information that builds on Stein (2008). In my model agents work to upgrade a product. In every period both agents make simultaneous decisions. Their decisions are in two stages. In the first stage, they decide whether or not to join the RJV. For the RJV to be formed in any period both agents must agree to join it. If the RJV is not formed, both agents work separately. In the second stage, agents simultaneously decide whether to work or not. While working is not an observable action, whether or not an agent upgraded the product is public information. Thus, if the agents are in an RJV, the failure of an agent to provide an upgrade can be the result of that agent shirking, or, bad luck.

The model has three key characteristics. First, actions are complementary: an agent in an RJV, assuming both agents work, have a higher chance of upgrading the existing product than if either works alone. Second, if the agents work together, they divide the monetary payoff from the upgrade, and earn less than if either had found the upgrade while working alone. Third, if at the end of any period an agent has not upgraded the product (either on his own or trough the RJV) then the game is over for that agent.

I first consider the case when abilities are observable and common knowledge. I show that there exists ability values for which being part of the RJV is a subgame perfect equilibrium

[38] Some of well-known examples of JV are Sony Ericsson (JV of Sony and Ericsson), Vevo (JV among Universal Music Group (UMG), Sony Music Entertainment (SME) and EMI), Hulu (created as JV between News Corporation and NBC Universal), Penske Truck Leasing (JV of Penske Corporation, Penske Automotive Group, and Mitsui Co), etc.

[39] https://www.investopedia.com/terms/j/jointventure.asp

[40] https://www.entrepreneur.com/article/236987

and that the first best outcome is attainable. I also show that there exist parameter values for which both agents want to cooperate but the RJV is not sustainable because one of the agents has an incentive to free-ride. In these cases, the first best outcome is not attainable. Note that, if the RJV is created, the main trade-off is between shirking to save the effort cost, and working to increase the joint probability of success. When the former is higher than the later, we can not have cooperative equilibrium.

I then consider what happens in the case of asymmetric information. In this situation, one player's ability (say player A - female) is known while the other player's ability (say player B - male) remains his private knowledge. However, it is known that player B's ability is either High or Low, with known prior distribution over possible types. Depending on player A's ability, there are four possible scenarios.

First, if player A's ability is too low she will free ride when she is matched with player B, regardless of whether B's ability is high or low. Therefore, no cooperation is attainable for any number of periods.

Second, if player A's ability is moderately low, she has an incentive to work if she is matched with a B player with low ability, but prefers to shirk if she is matched with high ability B player. Indeed, a high ability B player has a good chance of upgrading the product so player A prefers to let B work and save the effort cost. On the other hand, a low ability B player does not have a good chance of upgrading by himself so A is better off in expectation if she also works to find the upgrade. Thus, so long as player A believes that player B is a low type, cooperation is sustainable.

Third, if player A's ability is moderately high, neither she nor player B (whether high or low ability) have an incentive to free ride. In this case, cooperation will last until they both fail at the same time and exit the game simultaneously. This is because player A's ability is high enough to make him prefer helping B over shirking. On the other hand, player A's ability is not so high that low type player B has an incentive to rely on free riding.

Finally, if player A's ability is very high, low type player B has an incentive to shirk. A low ability B player prefers to save the working cost since A already has a good chance of upgrading the product, and B realizes that his own contribution would not be significant compared to his effort cost. A high ability B player on the other hand will have an incentive to work. Hence, player A agrees on cooperation with B as long as she believes B is a high ability player.

Since the dynamics of the first (never cooperate) and third (always cooperate) cases are obvious, I focus on the remaining two. These two cases (A has a moderately low or very high ability) can generate different dynamics for the RJV.

Let us first consider the dynamics of the RJV when A has a moderately low (observable) ability. In this situation, A will always prefer to be in the RJV. However, because she has an incentive to free-ride if she is partnered with a high ability player B and only works if she is partnered with a low ability player B, it is the B player who may want to leave the RJV.

The break up happens because player B accumulated high enough reputation so that player A has an incentive free-ride. In this situation, B prefers to break up and not let his partner shirk. There will be no re-establishment of RJV because the only way for player B to decrease his reputation is by failing to upgrade the product. However, since the agents are already working separately, failing in that period means that the game is over for the failed agent. I show that the best perfect Bayesian equilibrium is to join and work in the RJV while B's reputation is sufficiently low and break up (forever) once B's reputation reaches high levels. Also, I show that in this case there exists no equilibrium in which probability that RJV breaks up is zero. Therefore, in every equilibrium we will see break up with positive probability.

The second case I consider is the one in which player A has a very high (observable) ability. Let's assume that B is a low ability player. In this case, player B has a myopic incentive to free-ride. My model can generate cycles of working and free-riding, break-ups of the RJV and the re-establishment of the RJV. What leads to these cycles? When player B has low ability and is part of a RJV, he anticipates that A will eventually leave the RJV if he never works. B has a simple trad-off: work to try to increase his reputation, or shirk and guarantee that his reputation will go down. I found that there exists a mixed strategy perfect Bayesian equilibrium in which a low ability B player mixes between working and free-riding; and player A is willing to remain in the RJV so long as her belief that B is in fact a high ability player is high enough. The break-up and re-establishment of the RJV is the result of player A's threshold strategy. If she breaks up when B's reputation drops below the threshold, she should re-join the RJV if player B manages to increase his reputation back to the threshold level (which would happen if by chance B has more success than in expectation). These cycles can also happen if B is in fact a high ability player, but for the opposite reasons: even though B in this case works every period in the RJV, break-ups will happen if B is unlucky and A's belief regarding B's ability drops below the threshold. The RJV will be reestablished if B is no longer unlucky and is able to upgrade his product at the expected rate that a high ability player should.

The paper proceeds as follows: Section 2 introduces the theoretical model and the testable hypotheses we take to the laboratory. Section 3 describes the experimental design. Section 4 presents the experimental results. Section 5 analyzes the alternative theoretical model, and Section 6 concludes. Proofs, supplementary empirical findings, and full experimental instructions are presented in the Appendix.

3.1 Literature Review

In the next few paragraphs I will discuss related literature. The paper closest to mine is Stein (2008). The main departure I make relative to Stein (2008) is that I made the game simultaneous instead of sequential. Therefore, in my model a partner is just a complement who can increase success rates while in Stain (2008) a partner was necessary. My first result is similar to proposition 2 from Stein (2008) in which he states that there exists parameter values for which cooperation is sustainable. My model can be considered an extension of Stein (2008) model in the sense that I can allow for unobservable abilities, which makes it possible to generate different dynamics of RJV. Since in Stein (2008) project fails as soon as one of the player fail, they can not have a learning about the partner's ability. Which makes generating different dynamics of RJV imposible in his model.

My paper is also related to literature on strategic experimentation. Closest to my work is Bonatti and Hörner (2011) in which they examines moral hazard in teams over time. They found that free-riding leads not only to a reduction in effort, but also to procrastination. However, in their model partners have no option to be alone. Unlike them working together is the choice players make instead of endogenously given, so my paper can generate cycles of cooperation. Also, in all these models, a project is finished once the breakthrough is made while in my model a project is terminated when it fails.

My paper is linked to the game-theoretic literature on reputation effects started from Kreps and Wilson (1982) and Milgrom and Roberts (1982). The main result of this literature is that reputation effects can not harm the agent's payoff and typically it raises long-run payoffs, often to the agent's first-best. However, paper by Ely and Välimäki (2003) shows that the reputational concern of the long-run player to look good in the current period results in the loss of all surplus. This result is for the case when one of the agents is long lived and the other is short lived. They also show that if both agents are long lived then such losses can be avoided. I showed that reputation can have a harmful effect on cooperation even if both agents are long lived.

Grosskopf and Sarin (2010) investigate the impact of reputation in a laboratory experiment. They allow reputation to have either a beneficial or a harmful effect on the long-run player to test prediction of bad reputation from Ely and Välimäki (2003). They can not find support for a bad reputation and also beneficial effects are not as strong as theory suggests. I also conducted an experiment to test the prediction of my theory in the lab. I found that good reputation might have an adverse effects on cooperation.

The last paper that I want to mention in this section is Akcigit and Liu (2016). They investigate he role of information in innovation and showed that lack of collaboration has substantial cost for firms and for society.

3.2 Baseline Model

3.2.1 Setup

Consider the following situation. There are two players, A (female), and B (male), working on drug development. Time is discrete and denoted by $t \in \{1, 2, 3...\}$. Players are characterized by ability $p_i \in (0, 1)$ for $i \in \{A, B\}$, which is the probability of improving the existing drug. For the first part of the paper I assume that the player's abilities are common knowledge.

Every period is very similar to each other and players make at most two decisions. First, the players need to decide to be in the RJV or not. Second, players need to choose to work or not. In order RJV to be created **both** players have to decide to join it. Here is what happens if players stay alone or join the RJV:

- **Alone:** if a player is alone, she needs to decide to work or not. If she works, she incurs the cost of c and has p_i ($i \in \{A, B\}$) chance of improving the existing drag, which gives her utility of u and she moves to the next period. With probability $(1 - p_i)$ she fails to upgrade the drag and exits the game. If she decides not to work, she saves the cost but exits the game for sure because she did not improve the drag.

- **RJV:** if players are in the RJV, they still need to decide to work or not. If they both choose to exert costly effort, then they both incur the cost of c and have $(1 - (1 - p_A)(1 - p_B))$ probability of upgrading the drug and moving to the next period. In case of the upgrade, they both get only share from u (because upgrade was done in the RJV rathen than by individual). Let's denote share of i who is in the RJV with j by $\alpha(p_i, p_j) \in (0, 1)$, it is strictly increasing in i's own ability and strictly decreasing in her partner's ability. However, unlike the staying alone case, if one of the players choose not to work, then she still has a chance of upgrading the drug because the partner might have worked and upgraded the drug. Therefore, not working in case of RJV does not mean exiting the game for sure, and for the obvious reason, we call it free-riding (shirking).

Since every period is the same, we can say that the game for particular player continues until she has access to an upgraded drag from the previous period. That is, he continues if either one of this happens:

- she worked and improved the drug

- she was a member of the RJV, worked but failed to upgrade the drug, but her partner upgraded

- she was a member of the RJV, shirked, but her partner upgraded the drug

In other words, player exits the game and takes whatever she earned that far if at the end of any period she does not have access to a new upgrade, found either by her or by her partner. If player A exited the game then player B need to make only second decision from above (work or not), because option of being in the RJV is not available anymore.

Figure 22 summarizes the timeline described above. Assuming that both players are still in the game at the beginning of period t, then they first simultaneously decide to partner up or stay alone. After that, they need to choose to work or not (not working in the RJV means free-riding, while not working when a player is alone means exiting the game).

Figure 22: Timeline of the game assuming player are active at the beginning of period t

It is crucial to notice that working decision is not observable by partners. For example, if a player A does not provide the upgraded drug, player B cannot tell whether player A worked and failed or just shirked. However, if player A works and upgrades the drug then she can not hide it if she is in RJV. Therefore, if players are in RJV and one player upgrades the drug, then the credit for this update goes to both.

3.2.2 Payoffs

Payoffs are realized when the game is finished. This assumption does not play any role if the abilities are observable. However, in the next section, we assume players abilities are unobservable, and if we let the players see their earnings, then they can perfectly deduce their partner's ability from the sharing rule (explained above). Let's denote probability that at least one of the player finds the upgrade by $p_{ij} \equiv (1-(1-p_i)(1-p_j))$. Stage expected payoff matrix for player i is given in table 9:

	not work	work
alone	0	$p_i u - c$
RJV	$p_j \alpha(p_i, p_j) u$	$p_{ij} \alpha(p_i, p_j) u - c$

Table 8: Single period expected payoff matrix for player i with partner j

This table assumes that in the case of RJV, partner j works and not shirks. Let's look at the matrix. If player i is alone and is not working, then she does not incur any cost but also has no chance of improving the drug, so her expected payoff is 0. However, if she has a partner, assuming that partner works, she saves the cost and gets $\alpha(p_i, p_j)u$ with probability p_j, which is partners ability to upgrade the existing drug. If player i works alone, she incurs cost of working c and finds an upgrade with probability p_i, so in expectation she gets $p_i u - c$. Finally, if player i is partnered with player j, and they both work, then expected payoff from that round is $p_{ij}\alpha(p_i, p_j)u - c$, where p_{ij} is the probability that at least one of them find an upgrade for the drug.

From the expected payoff matrix we can say that player's total payoff depends on the number of upgrades she had with partner n_i^{RJV}, the number of upgrades she had alone n_i^{alone} and number of time she worked n_i^w. Therefore, we can write the total payoff U_i at the end of the game for player i as follows:

$$U_i = n_i^{RJV} * \alpha(p_i, p_j) * u + n_i^{alone} * u - n_i^w * c \tag{14}$$

The first term shows the total payoff while player i was in the RJV with player j. The second term is the payoff from the upgrades when a player worked alone, and the last term is the total cost of working from the entire game.

3.3 Observable Types

Let's look at the first best outcome of the game described above. To do so, assume there is a social planner who can dictate to the players what to do. The reason for looking at the social planner's optimum is to figure out when is creating RJV optimal for maximizing total output.

3.3.1 Social Planner

The main trade-off that the social planer face is having one good project versus having two mediocre projects. The intuition is that, if players collaborate, they will be able to find more upgrades for one project, then if any of them worked alone on the same project.

To solve for the first best outcome, we need to find expected payoff from the RJV and from working alone. Let's first look at expected payoff from the RJV given that both players work, equation 15.[41]

$$EU_{RJV} = \frac{(1 - (1 - p_A)(1 - p_B)) * u - 2c}{(1 - p_A)(1 - p_B)} \tag{15}$$

[41] All the derivations are shown in the appendix at the end of the paper.

If they both work separately, then the total expected payoff is just sum of expected utilities of each player, equation 16.

$$EU_{alone} = \frac{p_A * u - c}{1 - p_A} + \frac{p_B * u - c}{1 - p_B} \tag{16}$$

Finally, we just need to compare equations 15 and 16 to find ability pairs for which social planner prefers players to work in the RJV compared to working separately:

$$p_B \geq \frac{p_A * c}{p_A * u - c} \tag{17}$$

Equation 17 does not include optimal working decision. For completeness, we need to add two "working" conditions to get the social planner's optimum. For both player effort cost must be less than expected payoff from success, $p_A * u \geq c$ and $p_B * u \geq c$.

For graphical representation of condition 17, let's assume every upgrade of the drag is worth of $u = 10$, and the cost of working is $c = 2$. Using equation 17, the grey area on the figure 23 shows the ability values for which creating the RJV is socially optimal. For example if player A's ability is $p_A = 0.4$, then the social planer prefers her to work in the RJV with player B if $p_B = 0.8$ and prefers players to work separately if $p_B = 0.3$.

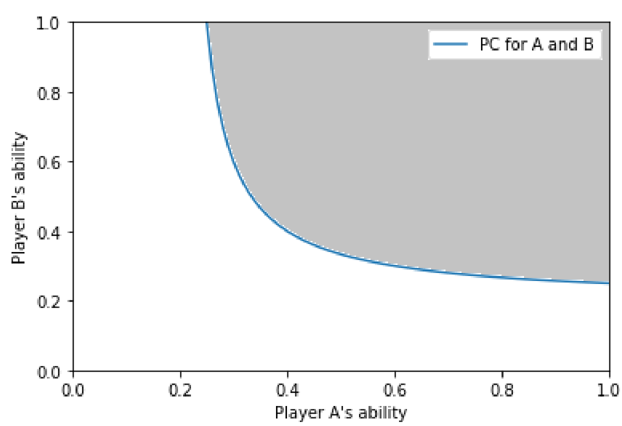

Figure 23: Social planner's optimum (the first best)

3.3.2 Sustaining the RJV

I define (cooperative) RJV equilibrium as a subgame perfect equilibrium in which both players want to be in the partnership and exert effort in every period. In this subsection, I explore the conditions under which the partnership equilibrium can be sustained. To do so, I hypothesize that such an equilibrium exists, and then check that at each date t, neither players have an incentive to deviate from the proposed equilibrium.

In the RJV equilibrium, each player prefers to be in the RJV and work in every period. Cooperation continues until both players fail simultaneously, at which point accumulated payoffs are realized, and the game ends.

From the definition of the RJV equilibrium, two conditions must be satisfied for both players: partnership constraint and effort constraint.

Partnership Constraint:

The timeline of the game shows that at the beginning of each period, first decision that the players make is whether to be in the RJV or alone. Assume it is period t and both players are still active (have not exited the game so far). Let's first look at player A's trad-off. She has to choose between joining the RJV and working alone. The expected payoff if she joins RJV, assuming that player B also joined and they both work is:

$$EU_A(RJV \ @ \ t) = U_A(t) + \frac{(1-(1-p_A)(1-p_B)) * \alpha(p_A, p_B) * u - c}{(1-p_A)(1-p_B)} \tag{18}$$

Where $U_A(t)$ is accumulated utility until period t. The second term on the right hand side of equation 18 is expected sum of utilities from future updates assuming that both players are in the RJV and work all the time.

Now let's look what would be player A's expected payoff if she deviates and works alone in period t (assuming that from period $t+1$ he will return to the RJV and stay there until the game ends):

$$\begin{aligned} &EU_A(alone \ @ \ t) = U_A(t) + p_A * u - c \\ &+ p_A * p_B \frac{(1-(1-p_A)(1-p_B)) * \alpha(p_A, p_B) * u - c}{(1-p_A)(1-p_B)} \\ &+ p_A * (1-p_B) \frac{p_A * u - c}{1-p_A} \end{aligned} \tag{19}$$

Equation 19 is easy to understand, the first term $U_A(t)$ shows accumulated utility until period t. The second terms show expected utility from the future updates if both players survives in period t. This happens with probability $p_A * p_B$ and starting from period $t+1$ they continue working in the RJV. Last term also shows expected utility from future updates but in case if player A updates and player B fails in period t. This happens with probability $p_A * (1-p_B)$ and player B exists the game. Therefore player A has to continue working on updates alone.

To find partnership constraint we can compare equations 18 and 19. Substituting $\alpha(p_i, p_j) = p_i/(p_i + p_j)$ instead of sharing rule, and solving inequality $EU_A(RJV\ @\ t) > EU_A(alone\ @\ t)$ gives us partnership constraint for player A:

$$p_B \geq \frac{p_A * c}{p_A * u - c} \tag{20}$$

Equation 20 shows that for every ability level of player A, she (player A) prefers to be in the RJV with player B if his (player B's) ability satisfies condition 20.

To get the partnership constraint for player B, we need to switch places of p_A and p_B in equation 20. Since everything is symmetric, we get the same condition for player B. Also, notice that this is the same condition as equation 17 from planners optimization. Which means that if social planner want players to work together in the RJV, then players agree and vice versa. However, this assumes that members of the RJV are working in every period. Therefore, before concluding that every socially optimal RJV will be created and sustained without a social planner, we need to take a look at the second constraint, namely effort constraint.

Effort Constraint:

To continue looking for the conditions under which the RJV equilibrium can be sustained, let's assume again that the players are still in the game and they both choose to join the RJV in period t. The next choice they have to make is whether to work or free-ride. The expected payoff from working is given in the equation 18 and the expected payoff from one shot deviation to free-ride in the current period is:

$$EU_A(free\text{-}ride\ @\ t) = U(t) + \frac{p_B(\alpha * u - c)}{(1 - p_A)(1 - p_B)} \tag{21}$$

Comparing expected payoffs from working (equation 18) and free-riding (equation 21), gives us the effort constraint for player A:

$$p_B \geq \frac{p_A^2 * u - p_A c}{c} \tag{22}$$

Similarly, we can find the effort constraint for player B. Let's again assume $u = 10$ and $c = 2$ for graphical representation of the effort constraints. Figure 24 shows effort constraints. Dotted curve corresponds to the effort constraint of player A. Which means that if for any level of player A's ability, player B's ability is above the dotted curve then, player A prefers

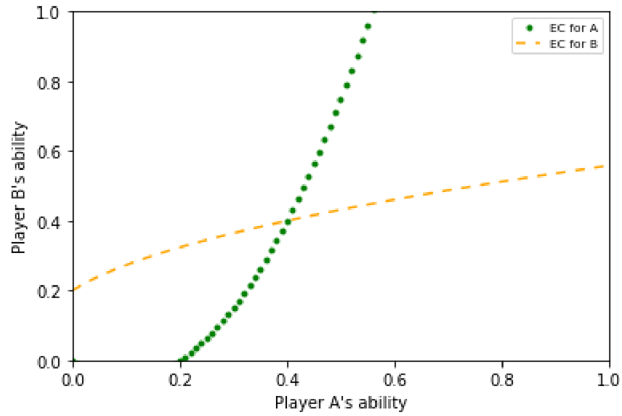

Figure 24: Effort constrains

to free-ride given that player B is working. For example, if player A's ability is 0.4, he will free-ride with the partner who is working and has the ability of 0.8 but chooses to work if B's ability is 0.2. The dashed curve represents player B's effort constraint.

The RJV Equilibrium:

We define the RJV equilibrium, as a subgame perfect equilibrium in which partnership and effort constraints are satisfied. We present both constraints above. Therefore we can join the arguments to find out when is maintaining the working RJV possible.

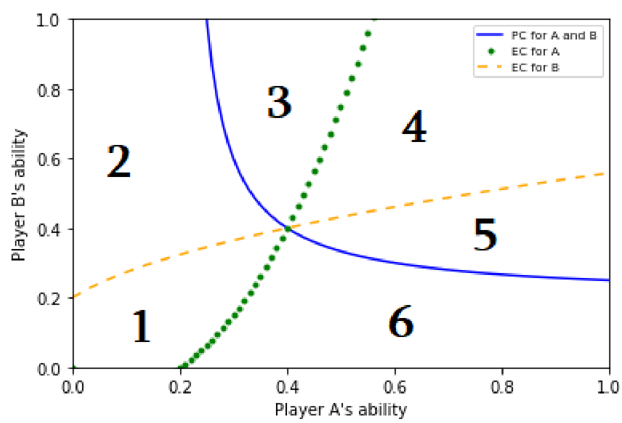

Figure 25: The RJV Equilibrium

All ability values can be divided into six groups as it is shown on figure 25. First, if ability pairs (p_A, p_B) do not satisfy the partnership condition (e.i. regions 1, 2 and 6) the RJV is not sustainable for obvious reasons. Let's look at region 3, from previous subsections we can say that ability values in that region satisfies partnership constraint but does not satisfy effort constraint for player A. Which means that even though both players want to be in the RJV, player A has a profitable deviation from working in the RJV to free-riding. Since abilities are observable, player B knows that player A will free-ride if they are in the RJV so he will not join the RJV. The same is valid for region 5 where arguments are reversed (now player B wants to free-ride). Finally, region 4 is such that both players satisfy the partnership and effort conditions. Therefore in the region 4 working RJV can be sustained.

From the previous section, we know that social planner prefers to make players work together if their abilities are in regions 3, 4, and 5. However, the first best outcome is sustainable only for the ability values in region 4. If the abilities are in the region 3 and 4, then social planners optimum is not sustainable because of the free-riding.

Proposition 1: If partnership constraints and effort constraints are satisfied for both players, then the RJV equilibrium is sustainable and it is the first best equilibrium in the sense that each player gets highest expected payoff compared to any other equilibrium.

3.4 Dynamics of RJV: Asymmetric Information

In this section, I assume only player A's ability is observable. Player B's ability is private information, and he can be either high type or low type with abilities, respectively p_B^H or p_B^L. Distribution over types is common knowledge. I define reputation as the probability of being a high type, denote by $\phi = \Pr(p_B = p_B^H)$. Player A updates his beliefs about player B after each round using Bayes rule. Depending on player A's ability, there are four different scenarios which results in four different dynamics of the RJV.

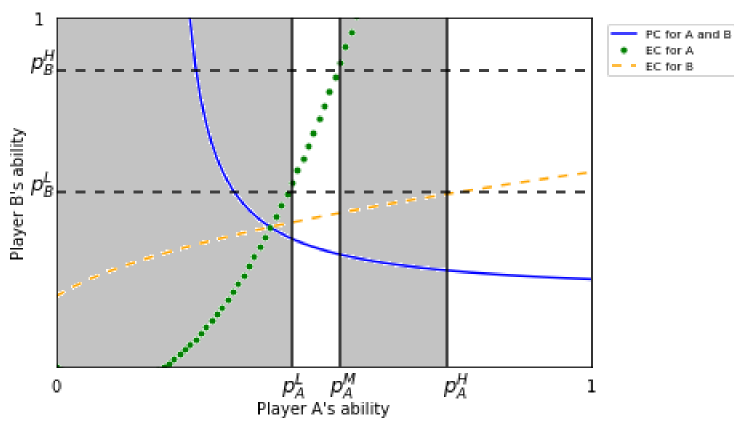

Figure 26: Four possible scenarios

All four cases are presented on figure 26. First, if player A's ability is lower than p_A^L, then creating RJV not possible. The reason is simple, either player B wants to be alone (if p_A is closer to 0) or player B wants to be in RJV with player A (if p_A is closer to P_A^L), but player A has an incentive to free ride, so player B refrains to join the RJV. Therefore in any case, if $p_A \leq p_A^L$, RJV will never be created.

Second, if player A's ability is between p_A^L and p_A^M, then she has an incentive to work in the RJV with a low type player B. However, if player B is a high type and A knows that, she prefers to free ride. Possible dynamics of second scenario is presented in subsection 3.4.1.

Third, if player A's ability is between p_A^M and p_A^H both players prefer to be in the RJV and work together until they fail. Joining and working in the RJV is weakly dominant strategy for both players.

Scenario four corresponds to the case when $p_A > p_A^H$. Player A's ability is too high, so low type player B prefers to be in the RJV but shirk. This is the opposite case of the second scenario in which player A had an incentive to free-ride. Subsection 3.4.2 discusses possible dynamics of this scenario.

Since the dynamics of the first (never cooperate) and third (always cooperate) cases are trivial, I focus on the remaining two. These two cases can generate different dynamics for the RJV.

3.4.1 Break up of the RJV

First we discuss possible dynamics of the second scenario from above. Let's say player A's ability is somewhere between p_A^L and p_A^M, denote it by p_A. In this section, I show negative consequences of a good reputation on RJV.[42] As I said player B could be either low ability p_B^L or high ability type p_B^H. Depending on B's actual type, player A chooses to free ride or work if they are in RJV together. The scenario that we are analyzing is presented on figure 27.

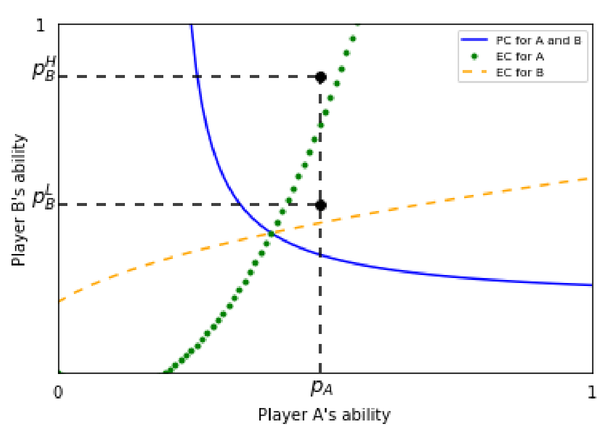

Figure 27: The effects of reputation on the partnership

[42]I define reputation as probability of being a high type, therefore good reputation is higher probability of being a high type.

As figure 27 shows if player A's ability is p_A, then p_B^H is above the EC for player A, therefor she prefers to shirk, given that player B is working in RJV. However, if player B is a low type with ability p_B^L (he is below EC for A), then A prefers to work rather then shirk in the RJV.

Intuition behind those incentives are as follows. If player B is a high ability type, and A knows that B is working in RJV, she prefers to save the cost of working and get whatever B generates, since B has a good chance of upgrading the drug. However, if B is a low type, player A prefers to work since player B alone in RJV can not upgrade the drug with a good enough chance. So, A prefers to work and increase chance of updating the drug rather than save the effort cost and let player B try to upgrade the drug alone in the RJV.

It is obvious that this scenario has multiple equilibria but we concentrate on the equilibrium that maximizes total expected output for both players.

Proposition 2: There exists the threshold reputation level ϕ^* such that, if initial reputation ϕ_0 is less than ϕ^*, it is the best perfect Bayesian equilibrium to work in the RJV every period until the reputation reaches ϕ^* and break up (forever) after that.

Proposition 2a: In every equilibrium of this game RJV breaks up with positive probability.

Proof of proposition 2: (not complete)

Only variable that changes period by period is the reputation of player B. Therefore we can write value function with the state variable ϕ_t:

$$V(\phi_t) = \max\{V_{alone}(\phi_t),\ V_{partner}(\phi_t)\} \tag{23}$$

Where:

$$V_{alone}(\phi_t) = \frac{p_i * u - c}{1 - p_i} \text{ for } i \in \{A, B\} \tag{24}$$

and

$$V_{partner}(\phi_t) = \max\{V_{work}(\phi_t),\ V_{shirk}(\phi_t)\} \tag{25}$$

We can use equation 25 to solve player A's effort constraint. From the previous section we know that if $\phi = 1$ then player A prefers to free-ride and if $\phi = 0$ he chooses to shirk. To solve equation 25 for player A for all $\phi_t \in (0, 1)$, let's first look each term separately:

$$V_{work}(\phi_t) = -c + \phi_t * [p_A(1-p_B^H)(\alpha_H * u + V(\phi_{t+1}^f)) + p_B^H(\alpha_H * u + V(\phi_{t+1}^s))] +$$
$$+(1-\phi_t) * [p_A(1-p_B^L)(\alpha_L * u + V(\phi_{t+1}^f)) + p_B^L(\alpha_L * u + V(\phi_{t+1}^s))] \quad (26)$$

$$V_{shirk}(\phi_t) = \phi_t * p_B^H(\alpha_H * u + V(\phi_{t+1}^s)) +$$
$$+(1-\phi_t) * p_B^L(\alpha_L * u + V(\phi_{t+1}^s)) \quad (27)$$

There is a heavy notation in equations 26 and 27, so before solving for effort constraint I will explain some notations. The share of player A when he is matched with high or low ability player B is denoted by $\alpha_\theta = \frac{p_A}{p_A + p_B^\theta}$ for $\theta \in \{H, L\}$. In period t player B either provides drug upgrade or not, in both cases player A updates his beliefs about player B's ability. Therefore, in period $t+1$ argument of the value function is either ϕ_{t+1}^s or ϕ_{t+1}^f, where s and f respectively denotes success and failure in upgrading the drug by player B.

Player A works if $V_{work}(\phi_t) \geq V_{shirk}(\phi_t)$. By comparing equation 26 and 27 we get that, player A works if:

$$\phi_t(1-p_B^H)(\alpha_H * u + V(\phi_{t+1}^f)) + (1-\phi_t)(1-p_B^L)(\alpha_L * u + V(\phi_{t+1}^f)) > \frac{c}{p_A} \quad (28)$$

If we show that LHS of equation 28 is decreasing in ϕ_t, we can say that there exists ϕ^* such that if $phi_t < phi^*$ player A prefers to work and shirk otherwise. This is because from the last section we know that if $phi_t = 1$ ($phi_t = 0$), then player A shirks (works). Therefor, we know at $phi_t = 0$ equation 28 holds and for $phi_t = 1$ equation 28 does not hold, since LHS is continues and increasing in ϕ_t, there must be ϕ^* such that equation 28 holds for equality (player A is indifferent between working and not working).

Why is the break up irreversible in this scenario? Since the reputation reached the threshold level, player B knows that if he join RJV again in the next period, player A will shirk. The same is true in any other periods if reputation is above the threshold. The only way for the reputation to decrease is if player B fails, but since players are alone, failure means that B exits the game, so they can not rejoin the RJV.

Q.E.D

3.4.2 Break up and re-establishment cycles of the RJV

In this subsection, I discuss possible dynamics of scenario four from the above. As I already mentioned scenario four corresponds to the case where player A's ability is above p_A^H from figure 26. Let's denote player A's ability by p_A which is above p_A^H. Therefore, if she is matched with low type player B and B knows that A is working, then B has one shot profitable deviation to free-ride. However, if B's a high type then both players want to work together in the RJV and none of them has an incentive to free-ride in any period. Figure 28 explains this scenario.

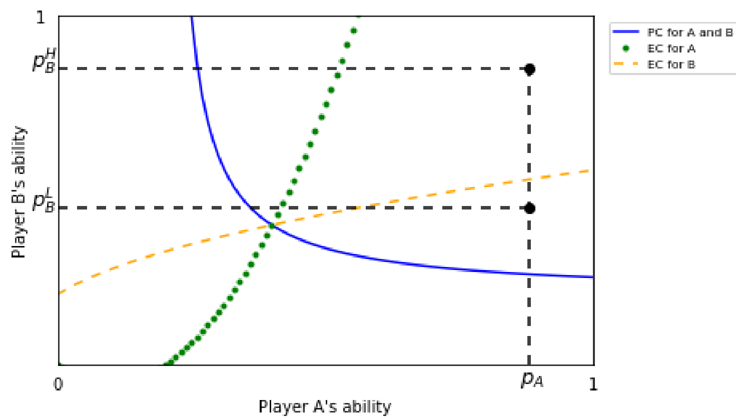

Figure 28: The effects of reputation on free-rider

Let me first explain rational behind players incentives. Incentive for high ability player B is obvious, he wants to be in the RJV and work in every period for two reasons. First, working increases his chance to have higher reputation, which is a necessary condition for player A to join the RJV. Second, he wants to work in every period because expected payoff from working is higher then from free-riding. For the low ability player B those two forces are working in the opposite direction. On the one hand he wants to work in order to increase his reputation and keep player A in the RJV. On the other hand he has myopic incentive to free-ride.

In the next proposition, I argue that if reputation is high enough, the partnership will be created, high type player B works in every period and low type player works in some periods. Even though low type player B has a myopic incentive to free-ride in every period, he sometimes works to build up the reputation and then cash it out by shirking for more periods. Once reputation drops below the threshold level, player start working alone but with positive probability they will re-join the partnership. So, we have break-up and re-establishment cycle on the equilibrium path.

Proposition 3: There exists perfect Bayesian equilibrium with threshold reputation level ϕ^* such that:

- If $\phi_t \geq \phi^*$:
  - Player A works in the RJV.
  - High ability player B also works in the RJV.
  - Low ability player B mixes between working and shirking in the RJV.
- If $\phi_t < \phi^*$:
  - Player A works alone
  - Both types of player B work alone to accumulate reputation and rejoin the partnership.

3.5 Experiment

The experiment was designed to study the effects of reputation building on collaboration. We use both within-subject and between-subjects comparisons to analyze changes in best response when potential partner has different ability.

All experiments were conducted at the Laboratory for Economics, Management and Auctions at Penn State University using oTree (Chen et al. (2016)). Subjects were recruited from the general undergraduate population, and each subject participated in only one session. A total of 64 subjects participated in our experimental sessions.

Study had two treatments and each consisted of four parts. The only difference between treatments were order of parts 1 and 2. The first three parts corresponds to game discussed in previous sections, let's call it a "Reputation Game". We asked simple lottery questions in Part 4 to assess subjects' attitude towards risk and concluded the experimental sessions by questionnaire.

3.5.1 Reputation Game

Let's first describe a typical game. A subject has fixed ability. As discussed in a theory part, ability corresponds subject's probability of success in a given round. A subject is paired with a computerized player of fixed ability. The game corresponds to what was described in the theory section, where player A is a subject and player B is a computerized player. In order to implement the possibly infinite nature of the game, I use a hybrid design where subjects make a sequence of decisions for rounds 1 through 4 and then I ask them to commit to a strategy in round 5. Round 1 through 4 all look the same. Figure 29 presents the timeline of each round of rounds 1, through 4.

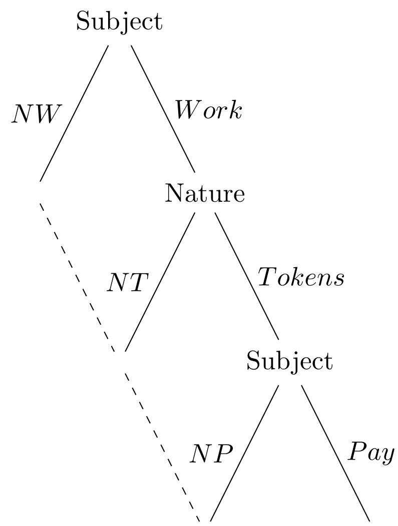

Figure 29: Timeline of each round of the reputation game.

As figure 29 shows, first decision that the subject makes is whether to *Work* or *Not Work*. As in theory, working is costly. If the subject decide to work, then nature moves. It gives the subject tokens with probability that corresponds her ability. If the subject finds tokens, she need to make last decision of a round, pay part of her tokens for computerized player to stay in the game or keep all her tokens and move to the next round without computerized player. If in any rounds, subject found token but decide not to pay, computerized player exits the game.

If subject decide not to work, or works but failed to find any tokens, then there might be two scenarios:

- Computer in: computerized player "Works" and find tokens for the subject with probability that corresponds computer's ability. If computer also fails then the game is over.
- Computer out: game ends and subject receives whatever she have earned that far.

Round 5 is a commitment round. If subject reaches round 5 and computer has not exited the game yet, she need to commit to either one of the strategies: work and pay the fee (if found tokens); work but do not pay the fee; or do not work (free-ride). If subject reaches round 5 but computerized player is not available anymore, then she commits to either working or not working strategies. Selected strategy in round 5 will be applied to every future rounds and the game ends in a round when player does not have an access to a new tokens that either her or computer found for her.

Now let's summarize game with parameters. Subject is paired with either low or high ability computerized player. If subject decide to work, she has 60% chance of finding 10 tokens and 40% chance of finding 0 tokens, in case of not work subject has 100% chance of finding 0 tokens. Work is costly and subject pays 2 tokens every time she decided to work. If subject worked and found 10 tokens, she need to decide to pay 5 or 7 token fee respectively for 60% and 95% ability computerized players to guarantee that the computer moves with her in the next round. If subject decides not to pay, she continues working alone for the rest of the game. If computer is still in the game and subject worked but failed or did not work then computer will try to help. Computerized player with low ability has 60% chance of finding 5 tokens and high ability has 95% chance of finding 3 tokens. If computerized player finds tokens, subject receives found tokens and moves to the next round with computer, if fails then the game ends.

3.5.2 Experiment Procedure

I conducted 4 experimental sessions during April 2019. Treatment 1 and treatment 2 each had 2 sessions, with number of subjects respectively 18, 10 in treatment 1 and 16, 20 in treatment 2. In total 64 subjects. Average (median) payment from experiment was $7.92 ($8.25), plus $7 participation fee. On average, sessions lasted about 1 hour. In the first three parts of experiment subjects played different version of the reputation games. In all those parts subject had an ability of 0.6. Difference between three parts was the ability of computerized player and information about it.

Let's describe treatment 1. In part 1 subjects played reputation game 15 times for practice, with low (0.6) ability computerized player. In part 2, subjects again played 15 times for practice with high (0.95) ability computer. In both part information about computer's ability was observable by subjects. In the third part, information about the computerized player's ability was not observable. However, subjects knew that it was computerized player either from part 1 or part 2 with equal chances each. Subjects played reputation game with unobservable computer type 20 times. Since subject experience both high and low ability computerized players in the first two parts, their goal in the third part was to identify type of the computerized player during first four rounds of each game and then to commit to their best strategy.

Treatment 2 was identical to treatment 1 with reversed order of the first two part. Subjects started with high ability computerized player, then moved to part two with a low ability computerized player. Part 3 was with computerized player either from part 1 or part 2 equal chances of each.

3.5.3 Theory predictions and Results

We can use section 3.4.1 to write down predictions for the reputation game. Notice that order of decision that subjects make is opposite then order of decisions from the model explained in theory part. This is not a problem for two reasons, first in the model with the reversed order, theoretical predictions for the parameter values that we chose are still the same, second, commitment round which is our primary interest has the same order then our theoretical model.

Computer's ability	Round 5 Predictions
60%	work and pay
95%	free-ride

Table 9: Theoretical prediction for commitment round

Parts 1 and 2 are for practice, so we do not have any theoretical predictions and we expect subjects to explore different options. Part 3 matters for payment and we can write down predictions. Subject's do not know the ability of computerized players, but they have 4 rounds to get information. As in theoretical part, we expect subjects to work and also pay computer fee in order to keep computer in the game and update their beliefs about computerized player's ability. If computer was successful in all 4 rounds (assuming subject is on the equilibrium path), then theory predicts that subject should commit to free-riding strategy as it is more likely that computerized player is a high ability. However, if computerized player fails even once, best response from subject is to commit to work and pay strategy as it is more likely that computerized player is a low ability.

We found support of the theory in our experiment. Figure 30 shows that subjects are twice more likely to commit to work and pay the fee strategy if computer failed at least once compared to never fail. Also players are 3 times more likely to commit to a free riding strategy if computer has a perfect strike compared to if it had at least one failure. As we see when computerized player had more success, therefore higher reputation, subjects decide to free ride instead of collaboration which means that reputation building has an adverse effects on collaboration.

3.6 Conclusion

I investigate possibility of collaboration among two firms who consider joining their resources to accomplish a common research goals. Even with a substantial support by governments, such units has a very high failure rate. In this paper, I first show that collaboration and the

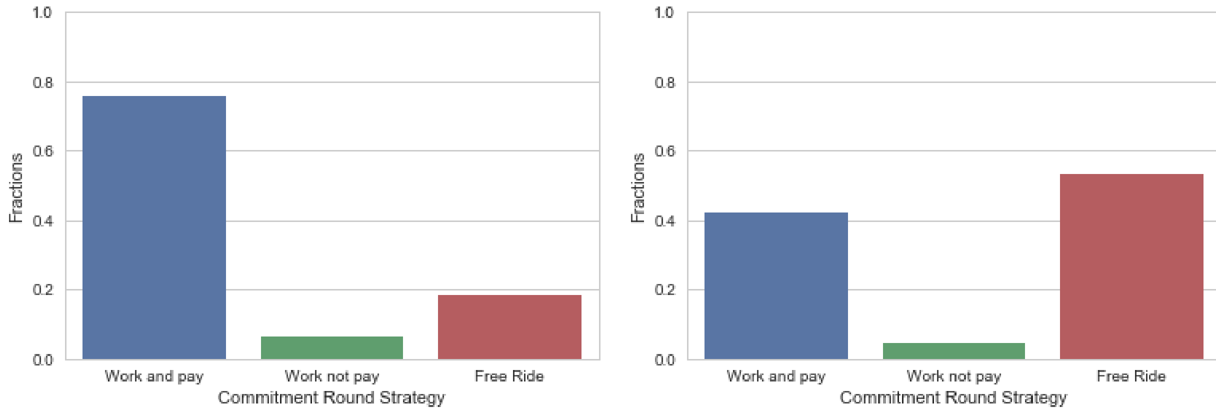

(a) Computer failed at least once

(b) Computer never failed

Figure 30: Strategies in commitment round

first best outcome is sustainable (achievable) for some parameter values when teammates abilities are observable.

If agents abilities are unibservable managing such teams becomes more challenging. I show that when there is asymmetric information about partner's ability, fear of free-riding might cause break up of collaborating unit and break up happens because of accumulated high reputation. Motivated by this finding I conduced a lab experiment where I show an adverse effects of reputation building on collaboration.

Last result of this paper is purely theoretical. I show that for some ability values we can have break up and make up cycles. In other words, firms start collaboration at first, but then break up and continue working separately because of depleted reputation. However, if agents manage to build their reputation by consecutive success, they might re-join their resources on the equilibrium path.

My findings can be applied by managers who consider joining their resources to collaborate. Also, very common in academia and generally in schools is creating teams for joint homework (papers). This process is challenging and random allocation might not be the one which gives as the first best outcome.

Appendix A

Appendix A Theoretical Discussion

What would explain the patterns of behavior that we find in the data? Can reference dependence, fatigue, gambler's fallacy, or hot hand fallacy explain the observed patterns? Let us start with reference dependence. One's personal best rating could act as a reference point for a player: a player ends a session whenever he sets his new personal best rating, and he plays longer otherwise. Reference dependence could only predict one type of behavior – stopping the session on a win. However, loss-stoppers can not be explained by reference-dependence unless we assume players have goals to lose a certain number of games every time they play. While theoretically possible, we have neither intuitive nor statistical support for such targets.

Another theory that we discuss is fatigue. Suppose that as a player keeps playing games, he gets fatigued over time; hence, fatigue would lead to worsened play over time. While lower last game scores are observed among loss-stoppers, it is opposite for win-stoppers, who have a much higher score in the last game. Moreover, even for loss-stoppers, we do not find any support of the fact that players' performance worsens as they play more games during a session.[43]

The last two theories that we consider are the gambler's fallacy and the hot hand fallacy. We call them belief-based explanations because both of them are based on players' beliefs about their future performance. The first one – the gambler's fallacy – implies the maturity of chances, if something happens more frequently than normal during a given period, it will happen less frequently in the future. The gambler's fallacy suggests that if a player wins a few games in a row, then a player may think that it is less likely that he wins again, and therefore, stops on a win. Similar logic applies if a player has lost a few games in a row, then he may think that he is more likely to win, therefore assuming person likes winning more than losing, he should stay and not end a session on a loss. While gambler's fallacy could explain some patterns found among win-stoppers, it goes against the behavior of loss-stoppers.[44]

The second belief based theory is the hot hand fallacy. Some athletes (as well as their fans) believe that if they succeed several times in a row, they have a "hot hand", i.e. higher chances of success in the next attempt too. It implies that if a player believes in the hot

[43] We ordered games for every session and find that fraction of wins is not affected by the order number of the game. We also conducted logit estimation to see if winning probability is affected by the order number of the game. We find that order number is not a statistically significant variable.

[44] In addition, if a player believes in the gambler's fallacy, that player's stopping decision should depend not only on the last game but on the game previous to last (lag 1 game). We show in Section 1.4.3 that the effect of lag 1 game result is weak compared to the last game result.

hand, that player should keep playing in case of win and stop after lose since it indicates that he does not have a hot hand anymore. Such behavior can predict the last game result of the loss-stoppers, however it goes against the behavior of win-stoppers.

To sum up, all the existing theories can explain one part of the story but is silent about the other. Win-stoppers might be the players who believe in gamblers fallacy and loss-stoppers be the one's who believe in hot hand fallacy. It can be that the entire population is a mixture of people who believe in different theories or principles. In this paper, we do not take a stand about underline psychological forces that motivate such behavior. Instead, we propose a theory that explains the actions found in the data and accounts for heterogeneity. Therefore, our model can accommodate win-stoppers, loss-stoppers, and neutral types together in the same model.

Appendix A Robustness

Figure 31 presents the players' population decomposition by types as we vary τ from 0 to .2. We see that behavioral types are robust to changing the allowed tolerance.

(a) Behavioral Type Fractions (b) Model Type Fractions

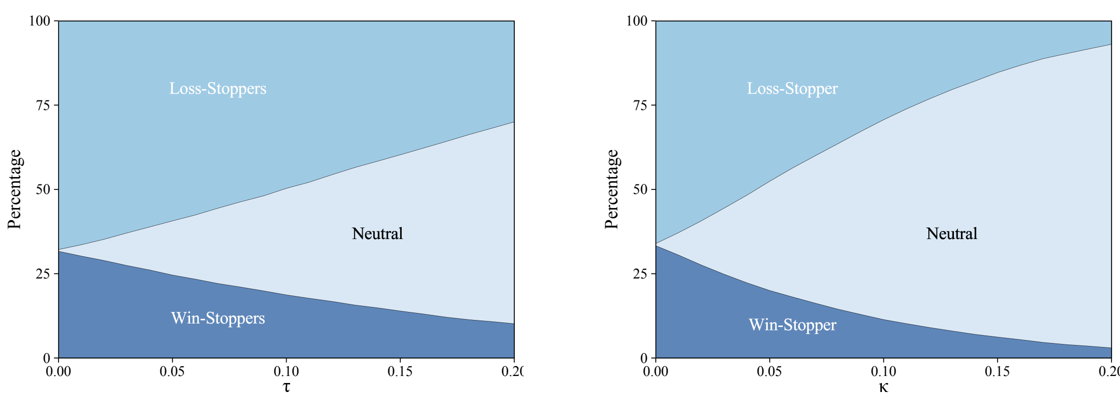

Figure 31: Type Decomposition

Similarly, changes in κ, also have consistent effect on model type decomposition, Figure 31b.

In this section, we examine whether a person's type changes over time and whether the rating in the chess game is correlated with the type classification. To study these questions we take two approaches: (i) we use existing data, and (ii) we collect additional data.

Figure 32a presents the break down of types over different rating categories. All three classification types are represented at every rating level. Moreover, the ratio of types for non-extreme rating levels are similar. Figure 32b presents the the distribution of the rating by type classification and we find that the distributions are practically the same. Kolmogorov-Smirnov test between win and loss-stoppers finds no difference with $p = 0.4$. Neutral types have on average lower rating than behavioral types with $p < 0.01$.

To further examine and look for any possible differences in rating and type classification we collect data for a special subset of chess players that have rating in top 2000. We do not find any qualitative difference in the distribution of behavioral type or in the ratio of win and loss-stoppers. These result suggest that the rating in the chess game is not related to the player's type.

(a) Number of Session (b) Session Length

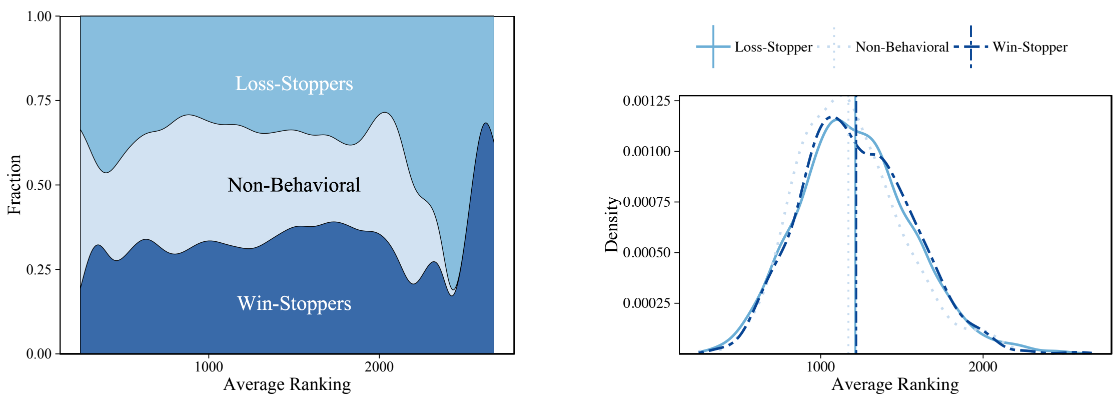

Figure 32: Ability and Behavioral Types

Our counterfactual shows that taking into account players behavioral type for matching algorithm can increase average number of games played during the session. One might think of several crowding out effects that increase of a session length might have. Even though we can not fully address such concerns without RCT, in this subsection we provide evidence that such effects is not likely.

The goal of counterfactual exercise was to increase number of games during the session, but the goal of market designer (platform) is to increase time spend on the platform. To address this issue we calculated correlation for every individual between minutes spend on the platform during a session and number of games played in the same session. Figure 33 shows that correlation between these two variables is high. Median correlation between minutes and games during the session is 0.98 across players.

Second issue that one might worry about is that asymmetric matching can cause fast games, in the sense that strong players can win with a weak players very fast. To show that that will not be an issue we calculated correlation between rating difference and minutes spend on a game. Rating difference gives as measure of how much better one player is

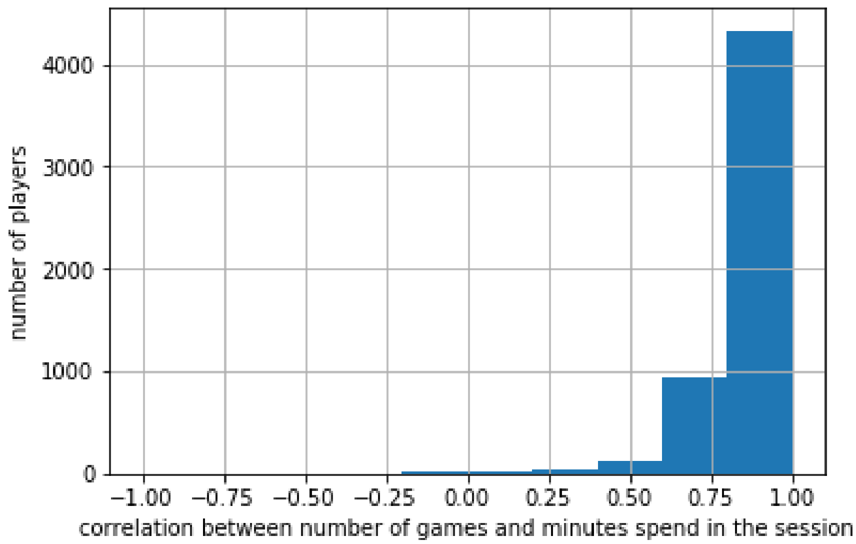

Figure 33: Correlation between minutes spend for a session and number of games

compared to another. Figure 34 shows correlation between how much better an opponent is and how much time the game lasts is close to zero for most of the players.

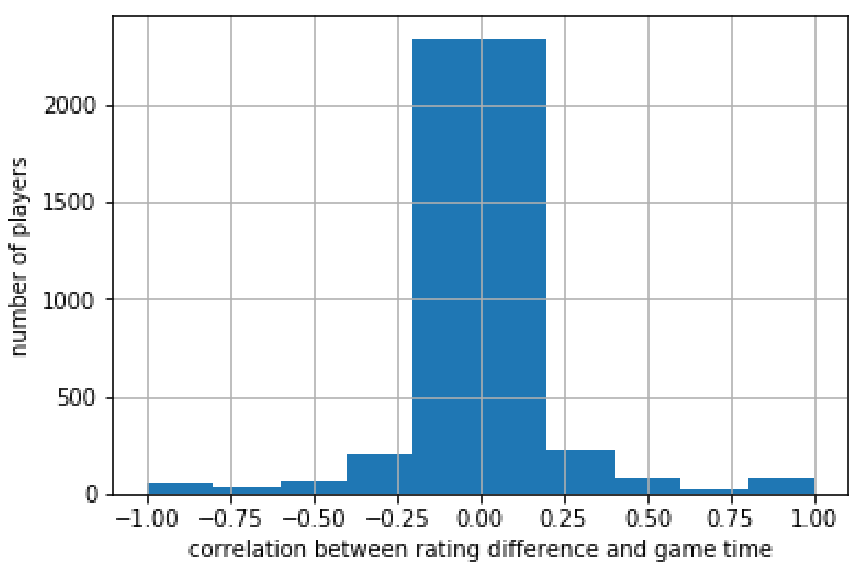

Figure 34: Correlation between minutes spend for a game and rating difference

One might worry that if a first session time increases during the day, it can decrease next session length (if players set out a certain amount of time to spend on the app every day). We find that correlation between number of sessions played during a day and average length of a session is 0.0002. We also find that one session length does not have any explanatory power on the length of the next session.

The last issue that we address here is about changing the type of the game. It might be that person who started a session with 5-minute blitz game can play shorter last game (for example 3-minute game) because he has only certain time allocated to play on the platform. If that is the case we should see that people change game types during the session. We find that 96% of sessions are homogeneous in the sense of the game type. This includes not only change if game type in the last game but during any other time. Which makes our argument even stronger that players do not choose last game type based on the remaining time that they have allocated for plying chess.

Appendix A Estimating CPH Model for Each Individual Separately

In our data set, an average player plays 348 session a year and an average session lasts 4 games. Therefore, we have sufficient data to estimate CPH model for each individual player. By examining the individual estimates we can identify win-stopper, loss-stopper and neutral types by looking at individual coefficients and the corresponding p-values. Since we already identified types using two different methods (one comparing winning fractions in the last and non-last games and other from the structural model), this is another robustness check of consistency of types over different methods.

Figure 35 presents the distribution of individual coefficients for the last game outcome variable. The figure highlights heterogeneity among the players.

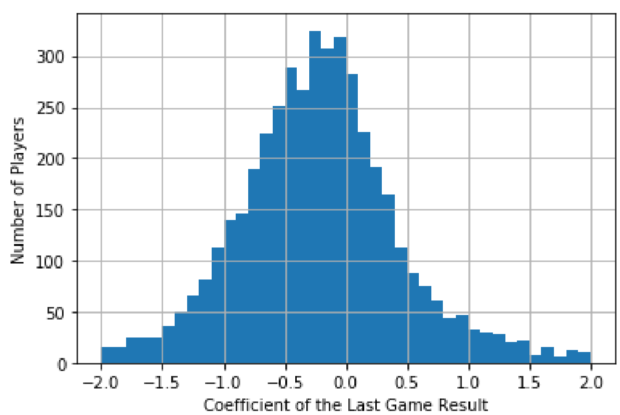

Figure 35: Individual coefficient distribution

Let us look at the differences in the hazard rate coefficient distributing of CPH model by behavioral types using the structural model estimation. Figure 36 shows the coefficient distributions for loss-discourage and win-stopper type separately. As we see win-stopper types have positive coefficients (Figure 36b), which means that win in the last game increases the chance that the player will end a session. For loss-stopper types (Figure 36a), the effect is opposite, win in the last game decreases the hazard of ending a session.

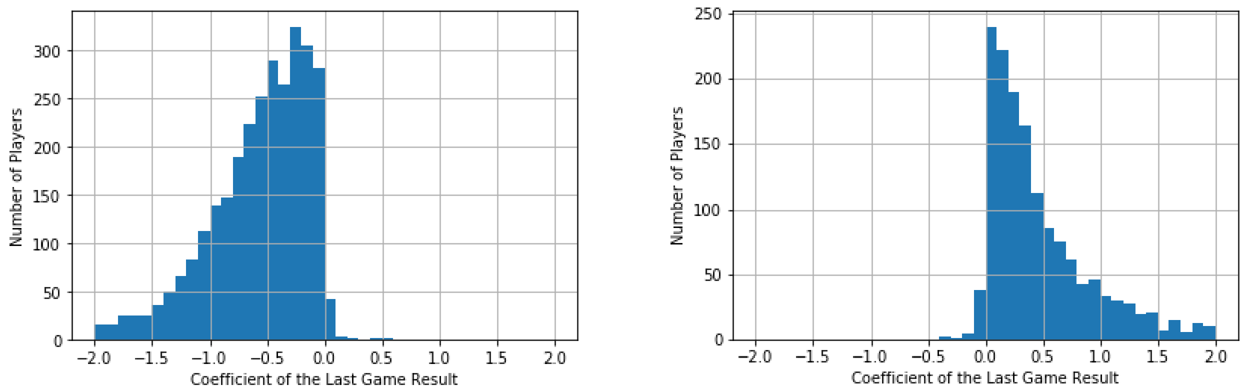

(a) loss-stopper

(b) win-stopper

Figure 36: Individual coefficient distributions for each behavioral type

Appendix B

Appendix B Derivations

***Claims 1-a, 1-b*:** First notice that if $p_0 = \mu_0$, $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, then $\mu_{1i} = p_{1i}$ for all $s_i \in \mathcal{S}$. Concavity of utility function reveals second cash out decision for each signal:

$$u\Big(\frac{x}{p_0} * p_{1i}\Big) > \mu_{1i} * u\Big(\frac{x}{p_0}\Big) \text{ for all } s_i \in \mathcal{S} \tag{29}$$

In other words, agent always cashes out i.e. $\mathcal{NCO} = \emptyset$. Now, we show that in this case agent is not investing. By contradiction assume she does:

$$u(x) \leq \sum_{i \in \mathcal{S}} \mu_{0i} * u\Big(\frac{x}{p_0} * p_{1i}\Big) = \sum_{i \in \mathcal{S}} \mu_{0i} \frac{\alpha_i^W}{\alpha_i^W} * u\Big(x * \frac{\alpha_i^W}{p_{0i}}\Big) < \sum_{i \in \mathcal{S}} \alpha_i^W * u(x) = u(x) \tag{30}$$

The first inequality is coming from equation 13. Second, equality is just simple algebra and definition of p_{1i}. Third, strict inequality is coming from the fact that $u(\cdot)$ is concave and lastly we know that $\sum_{i \in \mathcal{S}} \alpha_i^W = 1$. We get contradiction, so agent does not invests. We can do a very similar derivations for convex utility case, by reversing an inequality sign at equation 29 and 30.

***Claims 2-a, 2-b*:** First notice that if $s_z \in \mathcal{NCO}$ then from concavity of $u(\cdot)$ and from not cashing out condition:

$$u\Big(\frac{x}{p_0} * p_{1z}\Big) < \mu_{1z} * u\Big(\frac{x}{p_0}\Big)$$

we have that $\mu_{1z} > p_{1z}$. Which implies that $p_0 < \mu_0$, and since we assume $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, we can conclude that $\mu_{1i} > p_{1i}$ for all $s_i \in \mathcal{S}$.

Now we need to show the following implication:

$$\textbf{If } u\Big(\frac{x}{p_0} * p_{1j}\Big) > \mu_{1j} * u\Big(\frac{x}{p_0}\Big) \tag{31}$$

$$\textbf{then } u\Big(\frac{x}{p_0} * p_{1k}\Big) > \mu_{1k} * u\Big(\frac{x}{p_0}\Big) \tag{32}$$

for any $s_k \prec s_j$. By contradiction assume:

$$u\Big(\frac{x}{p_0} * p_{1k}\Big) \leq \mu_{1k} * u\Big(\frac{x}{p_0}\Big) \tag{33}$$

Substitute notations for $\mu_{1:}$'s and $p_{1:}$'s in equations 31 and 33 and simple algebra gives respectively:

$$\frac{\mu_{0j}}{\beta_j^W} u\Big(\frac{x}{p_0} * \frac{p_0 * \alpha_j^W}{p_{0j}}\Big) > \mu_0 * u\Big(\frac{x}{p_0}\Big) \tag{34}$$

$$\frac{\mu_{0k}}{\beta_k^W} u\Big(\frac{x}{p_0} * \frac{p_0 * \alpha_k^W}{p_{0j}}\Big) \leq \mu_0 * u\Big(\frac{x}{p_0}\Big) \tag{35}$$

Equations 34 and 35 implies:

$$\frac{\mu_{0j}}{\beta_j^W} u\Big(\frac{x}{p_0} * \frac{p_0 * \alpha_j^W}{p_{0j}}\Big) > \frac{\mu_{0k}}{\beta_k^W} u\Big(\frac{x}{p_0} * \frac{p_0 * \alpha_k^W}{p_{0k}}\Big) \tag{36}$$

Again by assumptions $\alpha_i^\omega = \beta_i^\omega$ for all $s_i \in \mathcal{S}$ and all $\omega \in \{W, L\}$, we can re-write equation 36:

$$\frac{u\Big(x * \frac{\alpha_j^W}{p_{0j}}\Big)}{\frac{\alpha_j^W}{p_{0j}}} > \frac{u\Big(x * \frac{\alpha_j^W}{p_{0j}}\Big)}{\frac{\alpha_k^W}{p_{0k}}} \tag{37}$$

Notice, $\frac{\alpha_j^W}{p_{0j}} > \frac{\alpha_k^W}{p_{0k}}$ and we get contradiction of equation 37.

We can use very similar arguments to prove the other side of claim 2-a and the entire claim 2-b for the convex utility function case.

Appendix B Graphs by treatment

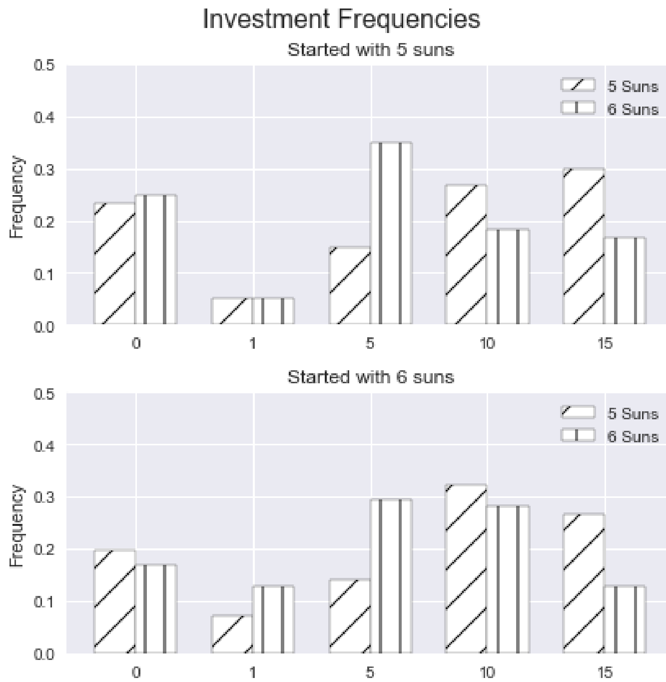

Figure 37: Investment frequencies by treatment.

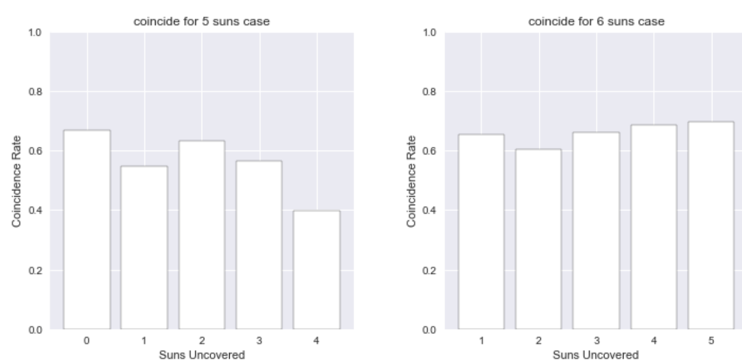

Figure 38: Matching the answers for the cash-out games and respective lottery questions separately for 5 and 6 suns cases.

References

Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, 101(2):470–92.

Aharoni, G. and Sarig, O. H. (2012). Hot hands and equilibrium. *Applied Economics*, 44(18):2309–2320.

Akcigit, U. and Liu, Q. (2016). The role of information in innovation and competition. *Journal of the European Economic Association*, 14(4):828–870.

Ambuehl, S. and Li, S. (2018). Belief updating and the demand for information. *Games and Economic Behavior*, 109:21–39.

Anderson, A. and Green, E. A. (2018). Personal bests as reference points. *Proceedings of the National Academy of Sciences*, 115(8):1772–1776.

Arkes, J. (2010). Revisiting the hot hand theory with free throw data in a multivariate framework. *Journal of Quantitative Analysis in Sports*, 6(1).

Arkes, J. (2013). Misses in "hot hand" research. *Journal of Sports Economics*, 14(4):401–410.

Avugos, S., Köppen, J., Czienskowski, U., Raab, M., and Bar-Eli, M. (2013). The "hot hand" reconsidered: A meta-analytic approach. *Psychology of Sport and Exercise*, 14(1):21–27.

Bonatti, A. and Hörner, J. (2011). Collaborating. *American Economic Review*, 101(2):632–63.

Brown, W. O. and Sauer, R. D. (1993). Does the basketball market believe in the hot hand? comment. *The American Economic Review*, 83(5):1377–1386.

Camerer, C., Babcock, L., Loewenstein, G., and Thaler, R. (1997). Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112(2):407–441.

Camerer, C. F. (1989). Does the basketball market believe in the'hot hand,'? *American Economic Review*, 79(5):1257–1261.

Cerulli-Harms, A., Goette, L., and Sprenger, C. (2019). Randomizing endowments: An experimental study of rational expectations and reference-dependent preferences. *American Economic Journal: Microeconomics*, 11(1):185–207.

Cervone, D., D'Amour, A., Bornn, L., and Goldsberry, K. (2014). Pointwise: Predicting points and valuing decisions in real time with nba optical tracking data, in '8th annual mit sloan sports analytics conference.

Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Coutts, A. (2019). Good news and bad news are still news: Experimental evidence on belief updating. *Experimental Economics*, 22(2):369–395.

Crawford, V. P. and Meng, J. (2011). New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101(5):1912–32.

Croson, R. and Sundali, J. (2005). The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of risk and uncertainty*, 30(3):195–209.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–38.

Elfenbein, D. W., Knott, A. M., and Croson, R. (2017). Equity stakes and exit: An experimental approach to decomposing exit delay. *Strategic Management Journal*, 38(2):278–299.

Ely, J. C. and Välimäki, J. (2003). Bad reputation. *The Quarterly Journal of Economics*, 118(3):785–814.

Farber, H. S. (2005). Is tomorrow another day? the labor supply of new york city cabdrivers. *Journal of political Economy*, 113(1):46–82.

Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *American Economic Review*, 98(3):1069–82.

Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, 130(4):1975–2026.

Frechette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109(8):2954–92.

Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, 17(3):295–314.

Green, B. and Zwiebel, J. (2017). The hot-hand fallacy: cognitive mistakes or equilibrium adjustments? evidence from major league baseball. *Management Science*, 64(11):5315–5348.

Grether, D. M. (1992). Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior & Organization*, 17(1):31–57.

Grosskopf, B. and Sarin, R. (2010). Is reputation good or bad? an experiment. *American Economic Review*, 100(5):2187–2204.

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

Imas, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *American Economic Review*, 106(8):2086–2109.

Kahneman, D. and Tversky, A. (1973). On the psychology of prediction. *Psychological review*, 80(4):237.

Koehler, J. J. and Conley, C. A. (2003). The "hot hand" myth in professional basketball. *Journal of sport and exercise psychology*, 25(2):253–259.

Kreps, D. M. and Wilson, R. (1982). Reputation and imperfect information. *Journal of economic theory*, 27(2):253–279.

Langer, E. J. (1975). The illusion of control. *Journal of personality and social psychology*, 32(2):311.

Milgrom, P. and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2):280–312.

Miller, J. B. and Sanjurjo, A. (2017). A visible (hot) hand? expert players bet on the hot hand and win. *Expert Players Bet on the Hot Hand and Win (December 15, 2017).*

Miller, J. B. and Sanjurjo, A. (2018). Surprised by the hot hand fallacy? a truth in the law of small numbers. *Econometrica*, 86(6):2019–2047.

Miller, J. B., Sanjurjo, A., et al. (2018). How experience confirms the gambler?s fallacy when sample size is neglected. *OSF Preprints*, 30.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2014). Managing self-confidence. *NBER Working paper.*

Morgul, E. F. and Ozbay, K. (2015). Revisiting labor supply of new york city taxi drivers: Empirical evidence from large-scale taxi data. In *Transportation Research Board 94th Annual Meeting*, volume 15.

Qadri, S. U. and Shabbir, M. (2014). An empirical study of overconfidence and illusion of control biases, impact on investor's decision making: an evidence from ise. *European Journal of Business and Management*, 6(14):38–44.

Rabin, M. and Vayanos, D. (2010). The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies*, 77(2):730–778.

Rao, J. (2009a). Inference of subjective sequences: When the gambler's fallacy becomes the hot hand fallacy. yahoo. Technical report, Research Labs, Mimeo.

Rao, J. M. (2009b). Experts' perceptions of autocorrelation: The hot hand fallacy among professional basketball players. *Unpublished technical manuscript. California. San Diego. Downloaded from http://www. justinmrao. com/playersbeliefs. pdf.(July 11th 2012).*

Rao, J. M. (2009c). When the gambler's fallacy becomes the hot hand fallacy: an experiment with experts and novices. *University of California, Working Paper.*

Rinott, Y. and Bar-Hillel, M. (2015). Comments on a'hot hand'paper by miller and sanjurjo (2015). *Available at SSRN 2642450.*

Sinkey, M. and Logan, T. (2014). Does the hot hand drive the market? evidence from college football betting markets. *Eastern Economic Journal*, 40(4):583–603.

Stein, J. C. (2008). Conversations among competitors. *American Economic Review*, 98(5):2150–62.

Stone, D. F. (2012). Measurement error and the hot hand. *The American Statistician*, 66(1):61–66.

Stone, D. F. and Arkes, J. (2018). March madness? underreaction to hot and cold hands in ncaa basketball. *Economic Inquiry*, 56(3):1724–1747.

Suetens, S., Galbo-Jørgensen, C. B., and Tyran, J.-R. (2016). Predicting lotto numbers: a natural experiment on the gambler's fallacy and the hot-hand fallacy. *Journal of the European Economic Association*, 14(3):584–607.

Sundali, J. and Croson, R. (2006). Biases in casino betting: The hot hand and the gambler?s fallacy. *Judgement and Decision Making*, 1(1):1.

Thakral, N. and Tô, L. T. (2017). Daily labor supply and adaptive reference points. Technical report, Working paper.

Tversky, A. and Gilovich, T. (1989). The cold facts about the "hot hand" in basketball. *Chance*, 2(1):16–21.

Wardrop, R. L. (1999). Statistical tests for the hot-hand in basketball in a controlled setting. *Unpublished manuscript*, 1:1–20.

Xu, J. and Harvey, N. (2014). Carry on winning: The gamblers' fallacy creates hot hand effects in online gambling. *Cognition*, 131(2):173–180.

Yaari, G. and Eisenmann, S. (2011). The hot (invisible?) hand: can time sequence patterns of success/failure in sports be modeled as repeated random independent trials? *PloS one*, 6(10):e24532.

VITA
Robizon khubulashvili

Education

Ph.D. in Economics, Pennsylvania State University, 2014-2020

M.A. in Economics (Honors), International School of Economics at TSU (ISET), 2011-2013

B.A. in Economics (Honors), Ivane Javakhishvili Tbilisi State University (TSU), 2007-2011

Working Experience

Summer 2019	Introduction to Microeconomics, PSU, Instructor.
Summer 2018	Money and Banking, PSU, Instructor.
2014–2019	Teaching Assistant.
2013-2014	BA Microeconomics, TSU, Instructor
2013-2014	AP Microeconomics, European School in Tbilisi, Instructor
Spring 2013	MA Game Theory, ISET, teaching assistant for Avner Shaked
Spring 2013	MA Microeconomics, ISET, teaching assistant for Florian Biermann
2013–2014	Researcher at ISET Policy Institute, Tbilisi, Georgia
Summer 2012	Researcher at ACF International, Tbilisi, Georgia

Professional Activities

2019 October	2019 North American ESA Meeting, Los Angeles, California
2019 July	2019 ESA World Meeting, Vancouver, Canada
2019 June	53rd Annual Conference of the CEA, Banff, Canada
2019 April	Pennsylvania Economic Theory Conference, Poster Session
2013 August	CERGE-EI Teaching Fellowship program, Prague, Czech Republic

Honors, Scholarships, and Fellowships

2013	Zurab Zhvania Prize for Outstanding Leadership, ISET
2011–2013	Future Leader Scholarship, ISET
2011–2013	Full scholarship for MA program, TSU
2008–2009	Presidential Scholarship, TSU
2007–2011	Full scholarship for BA program, TSU

Computer Programs

Python, R, Matlab, Stata, O-tree, Z-tree

Languages

Georgian (native), English (fluent), Russian (spoken)