Experimental Economics (2021) 24:87–117
https://doi.org/10.1007/s10683-020-09651-1

ORIGINAL PAPER

Is response time predictive of choice? An experimental study of threshold strategies

Andrew Schotter[1] · **Isabel Trevino**[2]

Received: 27 November 2017 / Revised: 12 February 2020 / Accepted: 27 February 2020 /
Published online: 12 March 2020

Abstract
This paper investigates the usefulness of non-choice data, namely response times, as a predictor of threshold behavior in a simple global game experiment. Our results indicate that the signals associated to the highest or second highest response time at the beginning of the experiment are both unbiased estimates of the threshold employed by subjects at the end of the experiment. This predictive ability is lost when we move to the third or higher response times. Moreover, the response time predictions are better than the equilibrium predictions of the game. They are also robust, in the sense that they characterize behavior in an "out-of-treatment" exercise where we use the strategy method to elicit thresholds.

Keywords Response time · Threshold strategies · Global games

JEL Classification C71 · C9 · D03 · D89

1 Introduction

There is a long tradition in psychology and neuroscience of using non-standard (non-choice) data to explain the process of choice. In these studies investigators look for correlates with neural measures inside the brain to explain observed choices in

✉ Andrew Schotter
andrew.schotter@nyu.edu

Isabel Trevino
itrevino@ucsd.edu

[1] New York University, New York, USA

[2] University of California, San Diego, San Diego, USA

an attempt to find a link between what people do and what they think. The exercise in many of these papers is to correlate observed behavior to neural measures with the arrow pointing from observable behavior to the brain.[1] In this paper we reverse this procedure and use internal brain processes to make point predictions about future choices. In particular, we use Response Time (RT) to predict observed thresholds in a global games experiment. Global games are binary action games with incomplete information where payoffs are determined by a state variable that is not known to agents. Instead, agents observe noisy private signals about the state variable and payoffs are determined in such a way that they find it optimal to take one action for realizations of their signal above some critical threshold, and another action for realizations below it. We define a subject's RT as the time that expires between when the subject first observes his signal and when he makes his binary choice. We can think of this as a subject's contemplation time.

Our results suggest that by looking at observations associated to the highest or second highest RT at the beginning of the experiment (first 25 rounds) we can predict observed thresholds (albeit not necessarily equilibrium thresholds) at the end of the experiment (last 25 rounds). In particular, the signals associated to the highest or second highest RT during the first 25 rounds of the experiment furnish estimates of observable thresholds at the end of the experiment that are both unbiased and comparable to those that could be made using the observable choices of subjects in the first half of the experiment.

Our aim in this paper is exploratory, in the sense that we are interested in discovering whether RTs can predict observed choices, whether there exists any interesting heterogeneity of types in the population that exhibit different RT patterns, and finally whether other existing models of RT (mainly formulated for single-agent choice situations) are useful in explaining our data.

In pursuing this agenda we use our RT results to characterize two types of subjects whom we call Intuitionists and Learners. Intuitionists seem to have an intuition about the use of a threshold strategy and even its precise value from the beginning of the experiment, but simply can't articulate what it is. Hence, when offered a signal that is below their threshold they quickly take one action and when offered one above it they quickly take the other action. However, when a signal close to their true threshold is received for the first time, Intuitionists spend a longer time contemplating it since it is not necessarily clear which action to take.

Learners, on the other hand, act as if they understand the structure of the game, and maybe even the benefits of a threshold strategy, but do not know the appropriate threshold to use. They learn their threshold through experience in the game and trial and error. As a result, Learners are more prone to make mistakes in initial rounds, in the sense that in early rounds they violate the dictates of their eventual future

[1] Such a process is counter to the typical revealed preference methods used by economists (see, for example, Gul and Pesendorfer 2008).

threshold. Once they converge on a threshold, however, their behavior becomes indistinguishable from that of Intuitionists.[2]

We present an "out-of-treatment" exercise where we analyze the results of a second experimental treatment where subjects play a global game but we use the strategy method and explicitly ask subjects to report thresholds to be able to observe their evolution over time. We find support for our characterization of Intuitionists and Learners by classifying 80% of the subjects in either one of these two groups.

This paper is by no means the first to use response times in economics.[3] For example, starting in 2006 and using a unique web site where a huge number of responses can be registered to play any listed game or decision problem, Ariel Rubinstein has been an early and persuasive advocate of the use of RTs.[4] For example, Rubinstein (2007) suggests that not all strategic choices are equivalent in the sense that some are "intuitive" and respond to some salient features of the games being played, some are "cognitive" and require more serious thought, while others are "reasonless" and appear random. What is interesting is that these differences can be seen in the RTs of the subjects with more cognitive choices taking more time to decide. In other words, those choices which by inspection of the game appear to be more sophisticated, when chosen, are the same choices that are associated with longer decision times. In Rubinstein (2008) RTs are used to separate subjects into fast and slow types and look to see how their decisions correlate across different decision problems. Again, the information provided by RTs is valuable in understanding the types of decision makers distributed throughout the population. Finally, in Rubinstein (2013) RTs are used to evaluate when a mistake has been made in a particular decision problem with the signature that mistakes involve lower RTs. Again, these results present evidence about the decision process that is hard to obtain by only observing choice data.

Other papers in the economics literature that use RT are Piovesan and Wengstrom (2009), who find a relationship between egoistic choice and RT in a dictator game where higher RTs are correlated to fairer outcomes, and Wilcox (1993), who measures RT as a proxy for decision cost in the laboratory to study the relationship between decision cost and incentives in environments with different levels of risk.

Our study differs from these papers since we study RT in order to make point predictions of choices, as opposed to correlating RTs to observed choices. In a similar spirit to our study, Chabris et al. (2009) study the allocation of time in individual decision making to elicit time preferences and find an inverse relationship between

[2] The distinction between Intuitionists and Learners that we are making is not very different from what happens when we ask two different people who was the director of a film. While one person may know the answer but is not able to recall it (it is "on the tip of my tongue"), the other may never have known it. If you mentioned names to the first type (the Intuitionist) she would be easily able to reject wrong answers because she would know the right answer when she hears it. The second type (the Learner) would have to go through a very different process and perhaps need to do a search of each name mentioned, since they know they never knew the answer.

[3] See Spiliopoulos and Ortmann (2014) for a discussion on the usefulness of RT in experimental economics.

[4] http://gametheory.tau.ac.il.

average RT and the difference in expected value between the payoffs associated to each of the possible choices. They find support to the optimization theory of Gabaix and Laibson (2005) and Gabaix et al. (2006) which predicts that agents will allocate more time to choices between options of similar expected utility than to choices between options of dissimilar utilities. However, Chabris et al. (2009) do not use RTs to make point predictions about choices. One paper that is close to ours in focus is Konovalov and Krajbich (2019) who look at response time to study indifference in preferences and who uses this indifference point to estimate parameters in a subject's utility function.

Other papers focus on RTs as an output of the Drift Diffusion Model (DDM), a model that has a long history in the neuroscience literature (see Ratcliff 1978; Busemeyer 1985; Ratcliff and McKoon 2008). For example, Clithero (2018) compares the predictions of a DDM that combines RT and choice data to the predictions of a logistic model of individual decision making and find support for the DDM approach. This paper is similar in spirit to our paper, but differs in two main respects. One is that the task in Clithero (2018) is an individual decision task (choosing between two food alternatives), while our paper studies a strategic environment. Second, Clithero (2018) uses RTs in combination with choice data to make predictions of future choices, while we look only at observations related to RT to make our predictions and then combine them with choice data to see the additional predictive power of RTs. Other papers studying the relevance of the DDM in economics include Fehr and Rangel (2011), Krajbich et al. (2014), Alos-Ferrer et al. (2016) , Fudenberg et al. (2018), Webb (2019), and Woodford (2014).

After presenting our results we explore the predictions in terms of RT of three alternative models for our experiment: the Drift Diffusion Model, the Directed Cognition Model, and the predictions in terms of RT that would emerge from the canonical global games model that we present in Sect. 2. We find mixed support for these models.

The paper is structured as follows. In Sect. 2 we present the model of global games used in the experiment. The experimental design is explained in Sect. 3 and our results are presented in Sect. 4. We discuss three alternative models in Sects. 5 and 6 concludes.

2 The global game

We use global games as our vehicle to investigate RT since the unique equilibrium in these games takes the form of a threshold strategy, which is easy for subjects to understand and execute in the lab. As we will see, signals close to the individual thresholds are harder to evaluate than those further away, and hence RTs should be reliable indicators of where individual thresholds lie.

A global game, as introduced by Carlsson and van Damme (1993), is a coordination game with incomplete information where payoffs depend on an unknown parameter θ, which we call the state of the world, and on the actions of other players.

Global games have been applied to a variety of economic situations such as currency crises, investment decisions, or political revolts.[5]

In this game there are two agents in the economy who have to decide whether to take action A or action B. Action B is a safe action and yields a payoff of zero in all states of the world. Action A is a risky action and taking this action has a cost of T. The payoff from choosing action A depends on the state of the world, θ, and on the actions of the other player. In particular, we can distinguish three different regions for the state θ that will determine how (and if) the action of the other player affects individual payoffs:[6]

- If $\theta \leq \underline{\theta}$, then action B dominates action A, regardless of the action of the other player:

$\theta \leq \underline{\theta}$	A	B
A	$-T, -T$	$-T, 0$
B	$0, -T$	$0, 0$

- If $\theta \in (\underline{\theta}, \overline{\theta})$, we are in the "coordination region" where action A yields a payoff of $\theta - T$ only if both players coordinate on this action. When only one of the players takes action A, his payoff is $-T$:

$\theta \in \left(\underline{\theta}, \overline{\theta}\right)$	A	B
A	$\theta - T, \theta - T$	$-T, 0$
B	$0, -T$	$0, 0$

- If $\theta \geq \overline{\theta}$, action A dominates action B, irrespective of the other player's action:

$\theta \geq \overline{\theta} > T$	A	B
A	$\theta - T, \theta - T$	$\theta - T, 0$
B	$0, \theta - T$	$0, 0$

In this game, however, players cannot observe the true value of θ, instead they receive noisy private signals about it. In particular, they know that θ is randomly drawn from a normal distribution with mean μ_θ and standard deviation of σ_θ, i.e.,

[5] See Morris and Shin (2003) for an overview on global games.

[6] In general, $\overline{\theta}$ and $\underline{\theta}$ are set in such a way that we can differentiate two dominance regions for θ (one for $\theta \leq \underline{\theta}$ and one for $\theta \geq \overline{\theta}$) and an intermediate region (for $\theta \in \left(\underline{\theta}, \overline{\theta}\right)$) which, in the presence of complete information, would exhibit multiple equilibria. Notice that in this intermediate region the optimality of taking action A heavily depends on the expectation that agents have about θ with respect to T. In order to make the game non-trivial, T is assumed to be strictly smaller than $\overline{\theta}$.

$$\theta \sim N\left(\mu_\theta, \sigma_\theta^2\right)$$

Once θ is realized, independent signals are privately drawn for each player according to a normal distribution with mean θ and standard deviation σ:

$$x_i \sim N(\theta, \sigma^2)$$

Given that players do not observe θ directly, once they observe their signal they base their decision to take action A or B on the expectations about θ and about the likely action of the other player. In particular, once they observe their signal players update their beliefs about θ and make inferences about the probability of θ being in either of these three regions. If they believe that θ might be in the intermediate region then players have to form an expectation of the likely action of the other player, since in this region players need to coordinate in action A in order for action A to yield a high payoff.

As first proven by Carlsson and van Damme (1993), in these type of games the information structure leads players to use a monotonic decision rule in which they take action B for low realizations of their signals, and they take action A for high realizations of their signals. This effectively means that agents use a threshold strategy such that they take action A if their signal is higher than a certain cutoff, $x^*(\sigma)$, and they take action B if their signal is lower than $x^*(\sigma)$. Formally, this decision rule can be written as:

$$a(x_i;\sigma) = \begin{cases} A & \text{if } x_i \geq x^*(\sigma) \\ B & \text{if } x_i < x^*(\sigma) \end{cases}$$

The threshold $x^*(\sigma)$ is defined as the value of the signal for which an agent is indifferent between taking action A or B.[7] This means that when an agent observes signal $x^*(\sigma)$, the expected payoff of taking action A is equal to the expected payoff of taking action B, which is zero in this case. Formally, if we assume that agents use threshold strategies in equilibrium, $x^*(\sigma)$ is the unique solution to the following equation:[8]

$$\begin{aligned} &E\left[\theta \mid x_i, x_j \geq x^*, \theta \in (\underline{\theta}, \bar{\theta})\right] \times \Pr(x_j \geq x^* \mid x_i, \theta \in (\underline{\theta}, \bar{\theta})) \times \Pr\left(\theta \in (\underline{\theta}, \bar{\theta}) | x_i\right) \\ &\quad + E\left[\theta \mid x_i, \theta \in [\bar{\theta}, \infty]\right] \times \Pr(\theta \in [\bar{\theta}, \infty] | x_i) - T = 0 \end{aligned} \tag{1}$$

We can see in Eq. 1 how expected payoffs depend on the value of θ and on the action of the other player. Recall that action A yields a payoff of $\theta - T$ under two

[7] Note that the value of the threshold depends on the precision of the signal, which in the case of a normally distributed signal is equal to the inverse of its variance. In this case, the precision of the private signals is equal to σ^{-2}.

[8] A unique solution to Eq. (1) is ensured as long as the private signals are precise enough with respect to the prior, i.e. when $\frac{\sigma}{\sigma_\theta} < K$, where $\sigma_\theta$is the standard deviation of the prior about θ and K is a constant that depends on the parameters of the model. This is a standard condition in the global games literature and it is met for the parameters used in the experiment (for a detailed discussion about the conditions for uniqueness see Theorem 1 in Szkup and Trevino 2019).

conditions. Either $\theta \geq \bar{\theta}$, or $\theta \in \left(\underline{\theta}, \overline{\theta}\right)$ and the other player also takes action A. The first condition is captured by the second term of Eq. 1 which simply corresponds to the conditional expectation of θ times the probability of θ being in this region. The second condition is captured by the first term of Eq. 1, which corresponds to the expected value of θ times the probability of coordinating with the other player (i.e. the probability that the other agent observes a signal $x_j \geq x^*(\sigma)$, which leads him to take action A as well), times the probability of being in the intermediate region $(\underline{\theta}, \bar{\theta})$, everything conditional of the private signal x_i. Taking action A always has a cost of T, irrespective of the value of θ, so we subtract it in Eq. 1. Finally we equate the expected value of taking action A to zero, which is the payoff of taking action B, to find the value of the threshold that equalizes the expected value of both actions. Note that, by definition, when an agent observes a signal that has exactly the same value as the threshold he does not have any strict preference over actions. This effectively means that for these signals agents are not sure about which action would yield a higher payoff in expectation.

In the RT analysis that follows we do not assume that the threshold used by subjects is the equilibrium threshold predicted by the theory. Since subjects do not necessarily use the equilibrium threshold they might have mistaken beliefs that make them indifferent between actions. Therefore, when we talk about high RTs being related to indifference between their binary choices, we do not necessarily refer to the theoretical equilibrium indifference portrayed in Eq. 1.

2.1 Parameters used in the experiment

The global games model presented above is governed by a set of parameters $\Theta = \left\{ \mu_\theta, \sigma_\theta, (\underline{\theta}, \overline{\theta}), T, \sigma \right\}$. For the experiment, the parameters chosen are the following:

$$\Theta = \{50, 50, (0, 100), 18, 1\}$$

In particular:

- The fundamental θ is randomly drawn from a normal distribution with mean 50 and standard deviation of 50.
- The coordination region is for values of $\theta \in (\underline{\theta}, \bar{\theta}) = (0, 100)$.
- The cost of choosing action A is $T = 18$.
- The standard deviation of the private signals is $\sigma = 1$.

3 Experimental design

We present the results of an experiment to analyze the role that RT has on predicting choices in global games. The experiment was conducted at the Center for Experimental Social Science at New York University using the usual computerized recruiting procedures. Part of the data generated by these experiments is a subset of the much larger data set generated by Szkup and Trevino (2019), whose emphasis was on the strategic play of the subjects and not on their RTs. The experiment was programed in z-Tree (Fischbacher 2007).

All subjects were undergraduate students from New York University. Our experimental design is closely related to the work of Heinemann et al. (2004), who test the predictions of the global games model of Morris and Shin (1998) and find clear support in the data for the use of threshold strategies. However, Heinemann et al. (2004) do not analyze RTs.

In each session subjects play the game for 50 independent rounds. Our treatments vary according to the type of action choice (direct choice vs strategy method). In the treatment with direct choice of action subjects observe signals and then choose actions (as portrayed by the model in Sect. 2), while in the strategy method treatment we elicit thresholds in every period before the subjects observe their signal in order to study the explicit evolution of thresholds over time.[9]

Overall, we present the results of five sessions where we had a total of 104 participants. Table 1 summarizes our experimental design.

Subjects were randomly matched in pairs at the beginning of each session and stayed with the same partner for all rounds. Each session lasted approximately 60 min and subjects earned on average $20.

The state θ is randomly drawn at the beginning of each round according to a normal distribution with mean 50 and standard deviation of 50. Once θ is drawn, one private signal is independently drawn for each subject from a normal distribution whose mean corresponds to the chosen value of θ and with a standard deviation of 1. In order to minimize the noise in RT observations, at the beginning of each round, subjects have to click on a button to observe the signal that was generated for them, and then they have to choose an action, for the treatments with direct action choice.[10] The time between when the button was clicked to observe a signal and the moment when the choice was made is our measure of RT. For the strategy method treatment, before observing a signal, subjects have to report the threshold above which they would be willing to take action A and below which they would be willing to take action B.

After each round, each subject observed his own private signal, his choice of action, the realization of θ, how many people in his pair chose A, whether the outcome was favorable to A, and his individual payoff for the round.

The computer randomly selected five of the rounds played and subjects were paid the average of the payoffs obtained in those rounds, using the exchange rate of 3 tokens per 1 US dollar.

[9] The RT analysis is performed for the direct action choice treatment only. We use the strategy method treatment to provide out-of-sample evidence that is consistent with our RT analysis.

[10] Having a subject click on a button gives more certainty in terms of when a subject actually first sees the signal, reducing the noise for cases when they might be day dreaming.

4 Experimental results

We first present the analysis of the treatment with direct action choice to establish the results based on RT estimations and characterize subjects' behavior. We perform the data analysis by studying RT in the first 25 rounds, when subjects are getting acquainted with the game and deciding on a strategy, to predict observed thresholds in the last 25 rounds, once subjects have, presumably, converged to a stable behavior. We use our choice-based results of the last 25 rounds as the objective choice measures against which we compare our predictions based on RT.

We then move on to the results of the strategy method treatment to evaluate the robustness of our earlier characterizations by performing some out of treatment estimations.

4.1 Choice based estimations

Just as in Heinemann et al. (2004) we find that over 90% of our subjects use threshold strategies during the last 25 rounds of the experiment. We say that a subject's behavior is consistent with the use of threshold strategies if the subject uses either perfect or almost perfect thresholds. A perfect threshold is characterized by taking action *B* for low values of the signal and action *A* for high values of the signal, with exactly one switching point. This effectively means that the set of signals for which a subject takes action *A* and the set of signals for which he takes action *B* are disjoint. This type of behavior is illustrated in panel (a) of Fig. 1, which has the signals a subject receives on the horizontal axis and a binary value (0 for action *B*, 1 for action *A*) on the vertical axis. For almost-perfect thresholds, we allow these two sets to overlap for at most three observations. This means that subjects take action *B* for low signal values and action *A* for high signal values, but these two sets can intersect for at most three observations. Such behavior is portrayed in panel (b) of Fig. 1 where we fit a logistic function to the observed last-25 round data of a specific subject.

We observe the use of threshold strategies in 90.48% of our subjects in the DA treatment over the last 25 rounds of the experiment. Once we have identified the subjects who use threshold strategies, we estimate the threshold for each subject by taking the average between the highest value of the signal for which a subject chooses action *B* and the lowest value of the signal for which he chooses action *A* in the last 25 rounds. This number approximates the value of the signal for which a subject switches from taking one action to taking the other action, which is how we define a threshold.[11] We find a median estimated threshold of the group to be 21.46 with a standard deviation of 20.32.[12]

[11] For a complete characterization of choice-based measures for thresholds in a global game see Heinemann et al. (2004) and Szkup and Trevino (2019).

[12] Notice that on average, subjects do not seem to follow the equilibrium threshold predicted by the theory, which corresponds to 35.31. This is consistent with findings in the global games literature (see Heinemann et al. 2004, and Szkup and Trevino 2019). However, the purpose of this study is not to establish optimality of thresholds with respect to the theory, but to predict observed thresholds with RT.

Table 1 Experimental design

Treatment	# Sessions	# Subjects
Direct action choice (DA)	4	84
Strategy method (SM)	1	20

4.2 Response time estimations

We show in this section that during the first 25 rounds of the experiment, if we consider for each individual the signal for which he has the highest or second highest RT, then either of those signals is an unbiased predictor of the threshold that the subject employs in the last 25 rounds of the experiment. In other words, we use RTs in the first 25 rounds to predict the observed thresholds in the last 25 periods, once behavior has stabilized. We discard the first round because there is in general a lot of noise in the RT data (e.g., subjects are getting acquainted with the interface).

To make our case, we look at the difference between the signals associated with a subject's highest and second highest RT and that subject's eventual, last-25-round, threshold. If either of these signals are predictive of last 25-round thresholds, we would expect to see these differences distributed around zero.[13] The frequency distribution of these differences, for the highest, second highest, and third highest RTs, together with their CDFs, are portrayed in Fig. 2 and summary statistics are presented in Table 2. As we can see, the medians of the distributions of differences between estimated thresholds and signals associated to the highest and second highest RT are not statistically different from zero at the 1% level of significance. For this reason, we can interpret the signals associated to the highest or second highest RTs as unbiased estimators of the observed thresholds in the sample. However, this is no longer the case for the signals associated to the third highest RT. The median of the differences between estimated thresholds and the signals associated to the third highest RT is different from zero to the 1% level of significance, and the distribution of these differences is statistically different from the distribution of differences corresponding to the highest and second highest RT to the 1% level of significance, using a Kolmogorov–Smirnov test.

To further investigate how robust our results are with respect to the predictions based on RT, we look also at the signals corresponding to the fourth highest RT, fifth highest RT, and so on. The graph in Fig. 3 presents the median of the distribution of differences between last-25 round thresholds and signals corresponding to the highest RT, second highest RT, third highest RT, fourth highest RT, up to the 24th highest RT, in the first 25 rounds of the experiment. In the horizontal axis we have the rank of the RT, starting from the highest at the origin to the 24th highest RT at the right end of the axis. On the vertical axis we have the median difference between the signals corresponding to the nth highest RT and estimated individual

[13] Note that we do not expect to observe these differences to be exactly zero because individual thresholds are estimated numbers and the probability of getting a signal realization exactly equal to this number is very small.

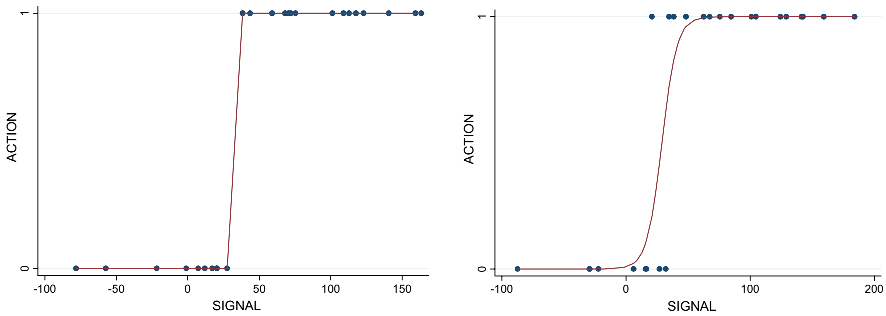

(a) Perfect threshold **(b)** Almost perfect threshold

Fig. 1 Examples of perfect and almost perfect thresholds

thresholds. Figure 3 illustrates how the signals associated to the RTs after the first or second highest are not good predictors of future choices, since the median difference of the signals associated with these higher RTs and the last-25 rounds estimated thresholds get further away from zero. Therefore, the accuracy of RTs as predictors of observed thresholds drops significantly when we move from the second longest to the third longest RT, and so on.[14]

In terms of length of RT, Fig. 4 shows the median RT for the highest RT, second highest RT, third highest RT, and so on. Notice that the median RT for the highest RT observations is 11.78 s and it is 7.64 s for the second highest RT, while the remaining RTs quickly decline and eventually converge to 2 s. Despite the fact that the median RT for the third highest RT is close to that of the second highest RT (6.46 s vs 7.64 s), they are different to the 5% level of significance (p value of 0.027).

It is relevant to point out that these are aggregate results. In other words, on average, across all subjects, either the first or second longest RT is an unbiased predictor of eventual thresholds. Obviously, as our figures indicate, there is a variance around these estimates. This implies heterogeneity in the sample, which we study later on.

In order to compare the predictive power of RT estimates to equilibrium predictions and other choice-based estimates, we refine our RT estimator. To aid us in this endeavor, we define the "Best Predicting Response Time" (BPRT) for each subject by looking at the signals associated with the highest and second highest RTs and selecting, for each individual, the signal that is closest to that subject's estimated threshold. Hence, for some subjects the BPRT will be associated to the signal with the highest RT, while for others it might be the signal with the second highest RT. This selection will facilitate the characterization of subjects into two different types. Table 3 recreates Table 2 for the BPRT

If the BPRT is meaningful to subjects, then we would expect their contemplation time to be different before and after they exhibit the BPRT. That is, we would expect subjects to spend less time thinking when they receive signals after their BPRT,

[14] Notice that the median difference between signal and threshold decreases again in period 8. This difference, however, is statistically different from 0 at the 1% level of significance.

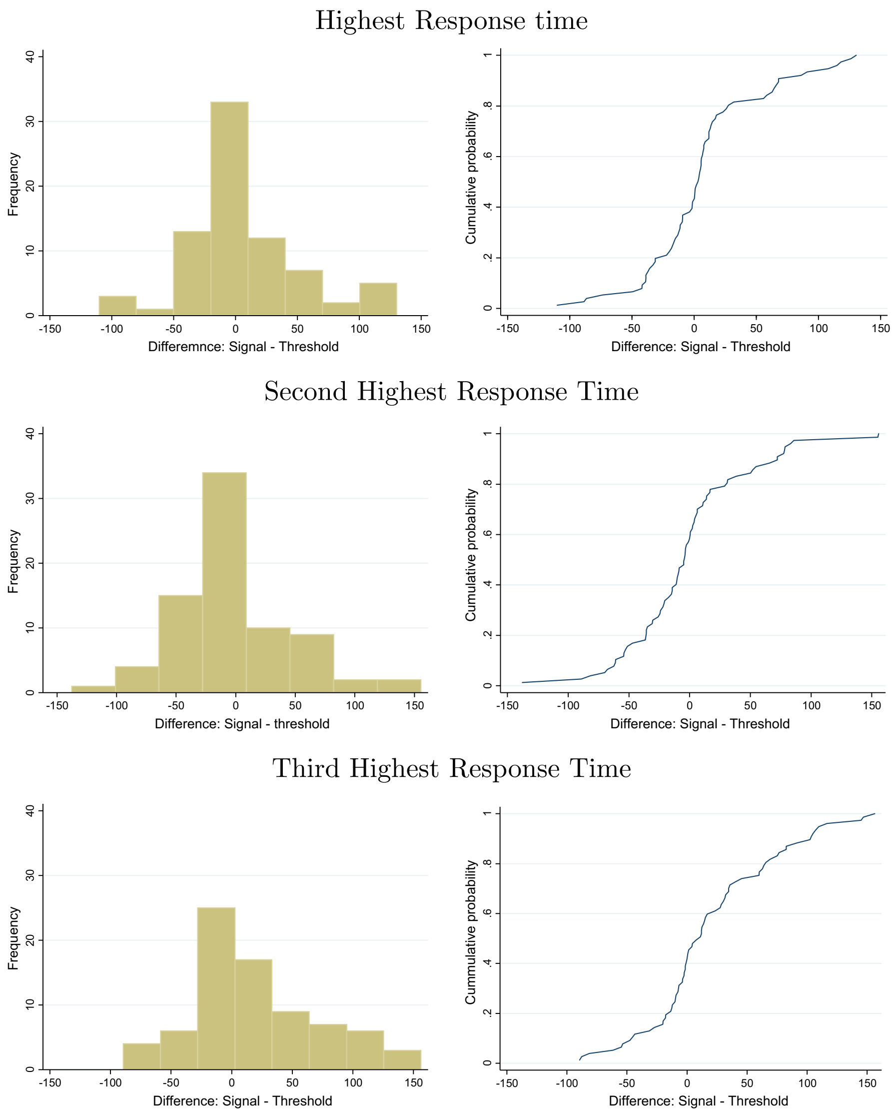

Fig. 2 Histogram and CDF of the difference between individual thresholds and signals associated to the highest and second highest RT, DA treatment

since experiencing their BPRT should make them feel more confident about what actions they should attach to each future signal, and hence they should spend less time thinking. Table 4 contains the median RT corresponding to the BPRT observations, observations before the BPRT, and observations after the BPRT. As is shown in this table, median RTs are lower after subjects have experienced their BPRT and this difference is significant at the 1% level. Note that when subjects receive the signal associated to their BPRT, the median subject spends 8.23 s thinking about

Table 2 Summary statistics of the difference between individual thresholds and signals corresponding to highest, second highest, and third highest RT, DA treatment

	Median (H0: $x = 0$)	Standard deviation
Highest RT	3.46	47.44
p value	0.431	
2nd highest RT	– 4.28	50.01
p value	0.326	
3nd highest RT	11.14***	53.65
p value	0.005	

Statistical significance to the 1% (***), 5% (**), 10% (*) level

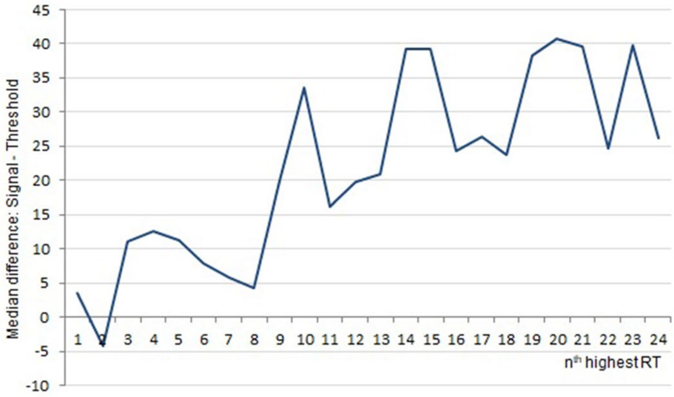

Fig. 3 Median differences between signals associated to RTs and thresholds, DA treatment

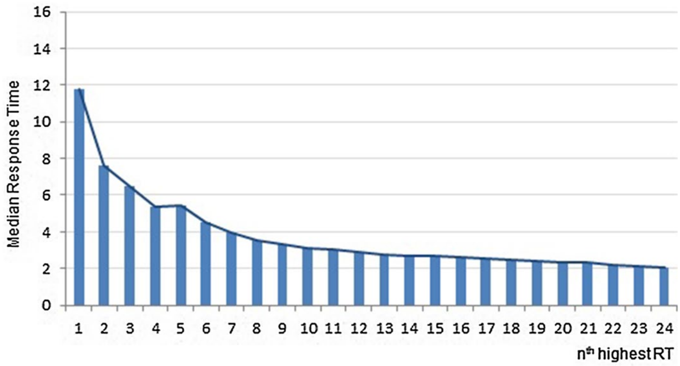

Fig. 4 Median RTs, DA treatment

it, while for signals received before their BPRT the median subject spends 3.07 s, and after the BPRT only 2.48 s. Therefore, we interpret the high contemplation time of the BPRT as reflecting the fact that, given the signal observed, it is not obvious

Table 3 Summary statistics of the difference between individual thresholds and signals corresponding to the BPRT, DA treatment

	Median (H0: $x = 0$)	Standard deviation
BPRT	− 0.54	26.96
p value	0.25	

Statistical significance to the 1% (***), 5% (**), 10% (*) level

Table 4 Summary statistics of distributions of RT before, after, and at BPRT, DA treatment

	Median	SD
RT exactly at BPRT	8.23	6.23
RT from $t \in \left[2, t_i^{BPRT}\right)$	3.07	3.31
RT from $t > t_i^{BPRT}$	2.48	2.04

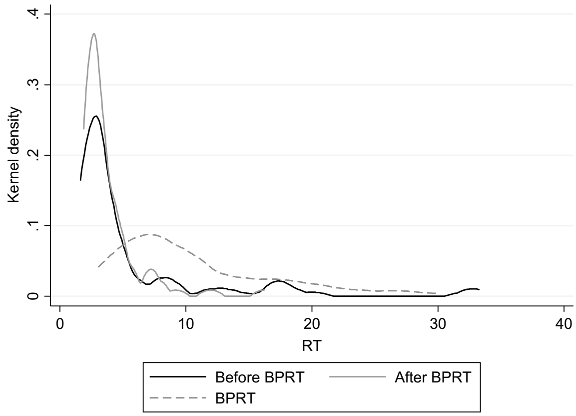

Fig. 5 Density of RT at BPRT and immediately before and after BPRT, DA treatment

for subjects what action to take, thus they need more time to deliberate. To provide further support for this argument, Fig. 5 plots the densities of observed RT at the BPRT and in the periods immediately before and after the BPRT. As we can see, the distribution of RT at the BPRT is strikingly different from the ones before and after the BPRT. Notice that the distributions immediately before and after the BPRT are similar to each other, but with the interesting difference that we observe a larger proportion of short RT immediately after the BPRT, supporting our findings.

The results presented so far give us an indication that subjects behave differently pre and post BPRT, which implies that their extended thinking at the BPRT imparts some knowledge. Our presumption is that this is a moment of insight where they gain some certainty about the value of their threshold and thus know what their behavior rule should be, so that after this discovery their RT is less sensitive to the signal they observe. This is consistent with the interpretation that a threshold strategy corresponds to the value of the signal for which a subject is indifferent between taking either action, that is, the signal for which the subject is not sure about which action to take. Moreover, the BPRT seems to serve as a structural break in terms of the behavior of RT. Table 5 contains the results of a random effects linear regression where the dependent variable is the logarithm of RT and the independent variables are Period, a dummy that takes the value of 1 for all periods up to the BPRT and zero otherwise (D(beforeBPRT) in Table 5), an interacted variable between this dummy and Period (Period*D(beforeBPRT) in Table 5), and a dummy that takes the value of 1 at exactly the BPRT period, and zero otherwise (D(BPRT) in Table 5). We can see a significant change in the way RT behave before and after the BPRT via the highly significant coefficients of D(beforeBPRT) and Period*D(beforeBPRT). Moreover, there seems to be a strong break point at the BPRT, as we can see from the high and significant coefficient of the BPRT dummy.

We now evaluate the predictive power of the signals associated with the BPRT to the equilibrium predictions of the theory. Table 6 shows the median difference between the equilibrium prediction (a threshold of 35.31) and the observed last-25 rounds thresholds. As we can see, the median difference is statistically different from zero to the 1% level of significance, indicating that the equilibrium predictions of the theory make rather poor predictions of actual eventual threshold behavior. This is in contrast to the signals associated to the BPRT, which, as Table 2 indicates, are unbiased predictors of future thresholds. Therefore, we can conclude that, on average, the signals associated to the BPRT are better predictors of observed thresholds than the theoretical equilibrium prediction.

With respect to our second comparison, we ask how the predictions of eventual thresholds made on the basis of our RTs (a non-choice variable) compare to what we could get if we measure thresholds in the first-25 rounds (choice based estimates). As we can see from Table 6 , the thresholds estimated for subjects in the first 25 rounds are also good predictors of future behavior, just as the highest or second highest RT (median differences are not statistically different to the 1% level of significance). This is interesting since it indicates that a non-choice variable (RT) can be as reliable a predictor of behavior as one based on choice data (first-25 round estimated thresholds). However, not all subjects that use threshold strategies in the last 25 rounds do so in the first 25 rounds. As we will see below, some of our subjects converge to a threshold by the end of the experiment by learning how to play the game, so their behavior in the first 25 rounds is not necessarily consistent with a threshold strategy.

4.3 RT and choice data: increasing predictive power

We use the BPRT to draw a better comparison between the predictive power of RT estimates and choice data. Figure 6 presents 2 histograms, analogous to those in Fig. 2, about the difference between the signals associated to the BPRT and the eventual threshold for each subject (panel a) and the difference between the first-25 round and last-25 round thresholds (panel b). These distributions are not

Table 5 Structural break of Log RT

Period	− 0.018***
	(0.002)
D(beforeBPRT)	0.405***
	(0.053)
Period*D(beforeBPRT)	− 0.034***
	(0.004)
D(BPRT)	0.94***
	(0.056)
Constant	1.387***
	(0.052)

Clustered (by subject) standard errors in parentheses

*Significant at 10%; ** significant at 5%; *** significant at 1%

Table 6 Summary statistics of the difference between individual thresholds and equilibrium predictions, DA treatment

	Median (H0: $x = 0$)	Standard deviation
Difference wrt theoretical equilibrium	13.86***	20.32
p value	0.0003	
Difference wrt first-25-round thresholds	0.062	19.99
p value	0.971	

statistically different from each other, which further suggests that RT estimates and choice-based estimates are equally as good predictors of future behavior.

Given this observation, we now ask whether RT analysis can provide additional predictive power over and above choice data. In other words, can these two estimates complement each other?

To answer this question we study, subject by subject, the relative importance of choice data and RT estimates for predicting future thresholds. In particular, we investigate whether we can think of a subject's last-25 round threshold as a convex combination of their first-25 round threshold (choice data) and the signal associated to their BPRT (RT prediction), or if some people's behavior is best predicted by just one or the other. We find, for each subject, the coefficient λ that best fits the equation below:

$$\begin{aligned} \textit{Last-25 round threshold} =& \lambda(\textit{First-25 round threshold}) \\ &+ (1-\lambda)(\textit{BPRT signal}) \end{aligned}$$

As we can see in Fig. 7, there are 3 types of subjects in the data. The subjects portrayed in the leftmost vertical bar corresponding to $\lambda = 0$ (30.435% of subjects) are better approximated by RT estimates, subjects portrayed in the rightmost bar

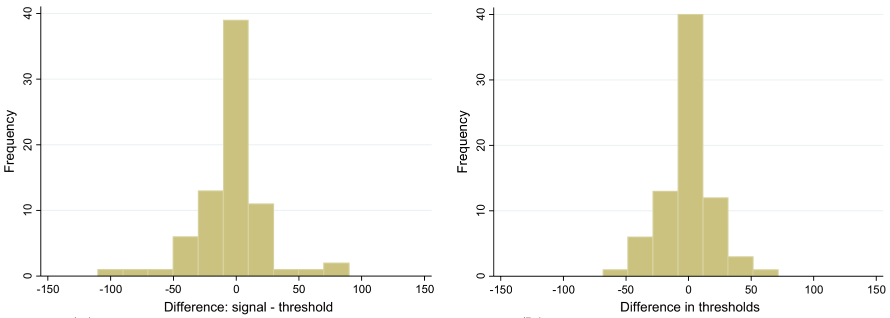

(a) BPRT signal - last-25 round threshold **(b)** first-25 round - last-25 round threshold

Fig. 6 Histograms of differences between last-25 round thresholds and **a** signal associated to BPRT, **b** first-25 round thresholds, DA treatment

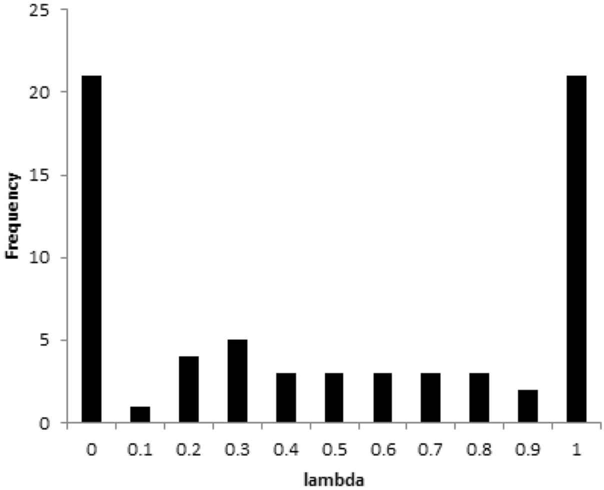

Fig. 7 Distribution of lambdas, DA treatment

corresponding to $\lambda = 1$ (30.435% of subjects) are better approximated by choice-data estimates, and the rest (39.13% of subjects) are a strict convex combination of both RT and choice-data estimates with parameter $\lambda \in (0, 1)$.[15] This implies that these two measures are actually complementary in predicting future choices because RT estimates approximate future thresholds for subjects whose first-25 round thresholds are not good predictors of their future behavior, and vice versa. That is, we see a clear increase in predictive power by using RT estimates in addition to choice data.

[15] The vertical lines at $\lambda = 0$ and $\lambda = 1$ include subjects whose threshold can be explained only by one of these measures because they're either greater or smaller than both measures. So these lines really correspond to $\lambda \leq 0$ or $\lambda \geq 1$, respectively.

4.4 Subject types

With these results in hand, we now investigate if our RT estimators can give us more information about the way in which subjects make their decisions. We make use of the BPRT to characterize the reasoning that leads subjects to think longer at that point by distinguishing between two different types of people, which we call Intuitionists and Learners. As described in the introduction, Intuitionists are subjects who act consistently with having an intuition about their threshold from the very beginning of the experiment. These subjects act as if they knew what their threshold should be, but can't fully articulate or recall it. However, once they observe a signal close enough to their eventual threshold, they stop searching. On the other hand, Learners are subjects who might understand what threshold behavior is, but who must learn from experience what their personal threshold should be.

This distinction between Intuitionists and Learners should be observable in the RT data.[16] While Intuitionists can be expected to ignore signals far from their implicit threshold and think hard when they observe a signal close to it for the first time, Learners may receive a signal close to their eventual threshold and ignore it, since they are learning about what their threshold should be and may not recognize a good signal when it first arrives. As a consequence, the first time an Intuitionist observes a signal close to her eventual threshold, that signal should become her BPRT, while a Learner may experience several such signals in early periods without those signals becoming a BPRT. We use precisely this distinction to classify subjects as Intuitionists if they do not observe a signal closer to their eventual threshold before they observe the signal associated to their BPRT, while we classify them as Learners if they do.

To give a simple example, say that a subject settles on a threshold of 25 in period 26-50 and in the beginning of the experiment receives signals 10, 55, 3, 22 (in that order). Say 22 becomes her BPRT, in the sense that she spends more time thinking about that signal than any other signal received in periods 1–25. Since there was no other signal closer to 25 received before that BPRT was determined, we will classify this subject as an Intuitionist. Now, say we have another subject who also settles down to a threshold in periods 26–50 of 25, but receives the following signals before receiving her BPRT: 10, 55, 3, 22, 24, 67, 29. Say 29 becomes this subject's BPRT. This means that she is spending a lot of time thinking about 29, despite the fact that she had earlier received signals closer to her eventual threshold. Since this subject did not deliberate when she received signals closer to her threshold than what eventually became her BPRT, she must have been learning when those signals arrived, so we would classify this subject as a Learner.[17]

[16] Our use of term Intuitionist differs from Rubinstein (2007). For Rubinstein intuitionists tend to have lower RTs to a given problem, while in our paper there is no particular difference between the duration of the first or second longest RTs for Intuitionists and Learners. What we find is that Intuitionists discover their threshold in earlier rounds than Learners.

[17] We have 37 intuitionists and 39 learners in our sample. 8 subjects were dropped from the sample because they do not use threshold strategies.

This distinction between Intuitionists and Learners manifests itself in the individual behavior of our subjects and the data they generate. More precisely, as a consequence of our classification, we expect certain differences in the behavior of subjects we classify as Learners and Intuitionists. For example, because Learners find the value of their threshold with experience while Intuitionists implicitly know it, Intuitionists should experience their BPRT in earlier periods than Learners (Learners need more time to learn, given identical signal distributions). Figure 8 illustrates this point by plotting, for each group, the distribution of periods corresponding to the BPRT, and Table 7, which presents the median of these distributions. Just as we expected, Intuitionists realize their BPRT in earlier periods than Learners. In particular, Table 7 indicates that while half of the Intuitionists experienced their BPRT by period 4, it took half of the Learners until period 14 to do so. Medians are statistically different at the 1% level.[18] As we see in Fig. 8 the two distributions of BPRTs appear considerably different, with the Intuitionists' distribution exhibiting far more of a right skew and a mass on earlier periods. These two distributions are statistically different to the 1% level using a Kolmogorov–Smirnov test.

In terms of stability of choices over time, we hypothesize that Intuitionists are more likely than Learners to behave in a manner consistent with their ultimate last-25 round thresholds from the very beginning of the experiment. In other words, if Intuitionists have a better understanding about their threshold while Learners need to learn it, then if we estimate two different thresholds -one for the first and one for the last 25 rounds- and then calculate the difference between them, we would expect smaller differences for Intuitionists than for Learners. This would imply that the thresholds of Intuitionists are relatively more stable over time than the thresholds of Learners, whose behavior may change due to learning.[19] To investigate this, we compute, for each subject, the difference (in absolute value) between their threshold in the last 25 rounds and their estimated threshold in the first 25 rounds. Figure 9 plots the CDFs of these differences, for Intuitionists and Learners separately. As we can see, the distribution of absolute value differences between first-25 and last-25 round thresholds for Learners stochastically dominates the one for Intuitionists. That is, the disparity between thresholds at the beginning and at the end of the experiment is larger for Learners than for Intuitionists. Table 8 reports, for each group, the median difference of first and last 25-round thresholds (in absolute value) and the standard deviation of the distribution of these differences, which further suggests more stable behavior for Intuitionists. For example, the median difference between first-25 and last-25 round thresholds (in absolute value) for Learners is 10.5, while

[18] Such a difference might lead one to think that perhaps a better way to classify subjects would be by calling them either fast or slow learners and choosing some arbitrary number of periods before a BPRT is determined to separate them. While this would have the benefit of being an exogenous classification scheme, it would offer no explanation as to why some subjects are fast and some slow and would be unable to offer any insights into the stability of behavior in our SM treatment to be discussed later on.

[19] It is important to note that the stable behavior that characterizes Intuitionists is not a consequence of matching. That is, it is not the case that Intuitionist subjects are better coordinated with their opponent in the game. In fact, looking at the composition of pairs Intuitionists are not more likely to be paired with other Intuitionists than Learners.

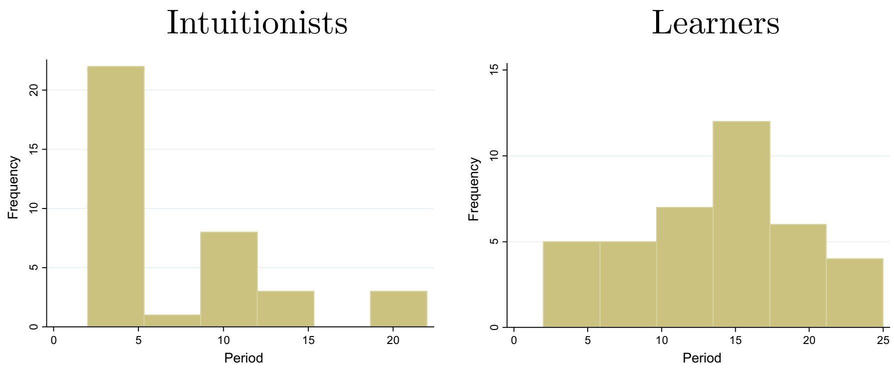

Fig. 8 Distribution of BPRT periods, by group, DA treatment

Table 7 Summary statistics of BPRT periods, by group

	Median	SD
Learners	14	6.32
Intuitionists	4	5.91

for Intuitionists it is only 5.73, and these medians are different to the 5% level of significance. A Wilcoxon test rejects the hypothesis that the sample of threshold differences in absolute value came from the same population at the 5% level. This implies that Intuitionists seem to have a clearer sense of what their threshold is in the earlier rounds of the experiment when compared to Learners, thus exhibiting more stable thresholds over time.[20]

As explained above, RT estimates and choice-based estimates from early rounds are not only similarly good predictors of future behavior, but actually complement each other. We recreate Fig. 7 but separating subjects by types in order to understand the relative predictive power of these two estimates for Intuitionists and Learners. This is illustrated in Fig. 10. For both types, approximately 39% of thresholds are better predicted by a strict convex combination of RT and choice data measures, similar to the overall sample. However, for the rest of the subjects, choice-data is better than RT estimates at predicting thresholds for intuitionists (38.89% vs 22%) and RT estimates alone are a better predictor of future thresholds for learners than

[20] One might wonder if Learners are just noisier decision makers than Intuitionists. To investigate this we look at average violations of individual thresholds in the last 25 rounds for these 2 groups by counting, for each subject, the number of times that their action in the last 25 rounds is not consistent with the threshold we estimate. We then take the average of these numbers for the subjects in each group. We find, on average, 2.11 violations for Intuitionists and 3.87 for Learners. However, the larger number of violations of Learners is mainly driven by 3 subjects. If we remove these outliers, we have, on average, 2.67 violations. Since these numbers are fairly similar, once we remove outliers, we remain agnostic about the potential role of RT in identifying noisy players in later rounds.

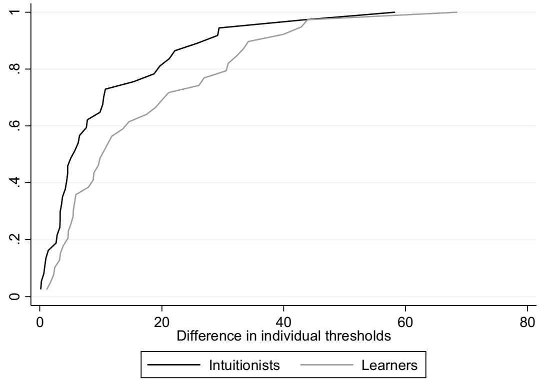

Fig. 9 CDF of differences of individual thresholds (in absolute value) from first 25 and last 25 rounds, by group, DA treatment

Table 8 Summary statistics of differences in individual thresholds from first-25 and last-25 rounds, by group, DA treatment

	Median	SD
Learners	10.5	15.23
Intuitionists	5.73	12.77

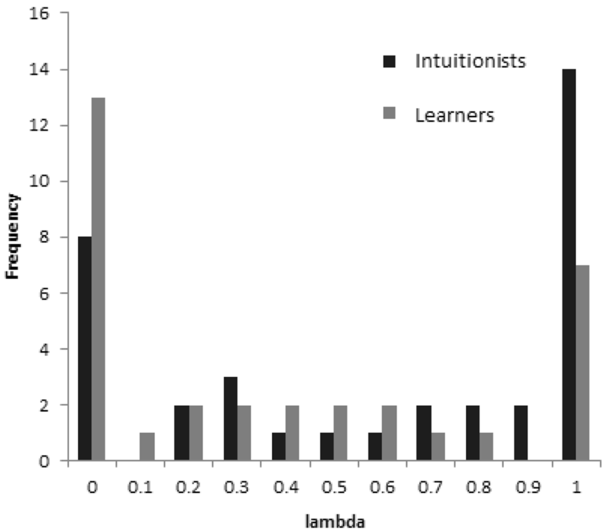

Fig. 10 Distribution of lambdas, by type, DA treatment

choice data (39.39% vs 21.21%).[21] These results give additional meaning to our characterization of subjects.

It is important to remember that in the last 25 rounds the subject behavior across types is indistinguishable from one another, since they all use thresholds and the mean estimated thresholds for each group are not statistically different from each other. This illustrates how RT analysis might give a broader insight about decision making than choice data alone. By studying RTs in the first rounds of the experiment we are able to distinguish how the different types of subjects come to realize their thresholds. In the following subsection we explore more explicitly the relative predictive power of RT estimates and choice-data.

4.5 Out-of-treatment predictions

One exercise that has proven to be informative (see Caplin and Dean 2015) is to axiomatize the behavior of economic agents and then characterize what the data from an experiment must look like if subjects behave in a manner consistent with those axioms. In other words, in a revealed preference type of exercise, one looks to see what the implications are for choice data of behavior that satisfied a set of assumptions or axioms. In this section of the paper we would like to ask a similar question with respect to our characterization of Intuitionists and Learners. While we have proposed no axioms, if our characterization of types is relevant we should be able to distinguish between subjects based on their exhibited behavior in the first rounds of the experiment.

The question is simple: If we performed a treatment where, instead of giving subjects a signal and then asking them to choose an action, we asked them before each period to state a threshold or cutoff level for realized signals above which A would be chosen, but below which B would be chosen, would we be able, from observing their reported thresholds, to classify subjects as Learners and Intuitionists? In other words, we are asking if we can observe Intuitionist and Learner behavior in an out-of-treatment exercise. To answer this question we make use of the data from our SM treatment, where subjects play the same game as before but where, instead of observing signals and choosing actions directly, we use the strategy method to ask subjects to report their threshold.

If our characterization of types is correct, then we should observe a group of subjects who report very stable thresholds from the initial rounds of the experiment, which would correspond to Intuitionists and another group, the Learners, whose

[21] It is not surprising to see that the thresholds of Intuitionists are better approximated by choice data since their behavior is in general very stable. To look at violations of individual thresholds in the last 25 round, for each group, we count, for each subject, the number of times that their action in the last 25 rounds is not consistent with the threshold we estimate. We then add these numbers and divide them by the total number of subjects in each group. We find, on average, 2.11 violations for intuitionists and 3.87 for learners. However, the larger number of violations of learners is mainly driven by 3 subjects. If we remove these outliers, we have, on average, 2.67 violations. Since these differences are not too stark (controlling for outliers), we remain agnostic about the role of RT in identifying noisy players in later rounds of the experiment.

period to period thresholds should exhibit far more variability in early rounds due to experimentation, and then stabilize. Our results indicate that by looking at the evolution of reported thresholds throughout the 50 rounds of the experiment we can categorize 80% of the subjects in this fashion. To perform this characterization we look at subjects who show some stability in reported thresholds in the last 25 rounds, to ensure convergence of behavior, and measure the standard deviation of their individual reported thresholds in the first 25 rounds. The data gives us a straight forward distinction between the subjects that we can potentially categorize as Learners and Intuitionists, in the sense that there is a group of subjects that exhibit low individual standard deviations of reported thresholds in the first 25 rounds (0–4.26) and a group of subjects with very large standard deviations (9.95–25.54). To give a better idea for how Intuitionists and Learners differ, consider Fig. 11 which offers an example of the evolution of reported thresholds for one subject categorized as Intuitionist (left panel) and one as Learner (right panel). Figures 2 and 3 in the online appendix plot these graphs for all subjects categorized in either of these groups.[22]

As we can see, the panel on the left offers a perfect picture of what our archetypal Intuitionist should look like in the SM Treatment. Previously we characterized Intuitionists as subjects who have a good idea of what their threshold should be but cannot explicitly verbalize it. In the DA treatment, subjects observe signals and implicitly set a threshold, whereas in this treatment (SM) they are forced to report it. This implies a different psychological process when establishing an action rule, which forces Intuitionists to explicitly verbalize their threshold. In line with our original description, Intuitionists act consistently with their threshold from the beginning of the experiment. Learners, on the other hand, are defined as subjects who are not quite sure what the right threshold should be and thus experiment in the initial rounds (see right panel in Fig. 11). For the SM treatment this would imply that Learners set many different thresholds in the initial periods, and then converge to a threshold. To support this claim, in Table 9 we present summary statistics for the individual period-to-period changes in reported thresholds in the first 25 rounds, by groups, and we find that the median period-to-period change in reported thresholds for Intuitionists is 0, while for Learners it is 6, with standard deviations of 2.51 and 16.36, respectively. Medians and standard deviations are significantly different at the 1% level. As a result, it seems clear that for one group of subjects (whom we label as Intuitionists) there is very little variability in the thresholds they set in early rounds while in another (whom we label Learners) there is quite a bit of variation. This is in line with our categorization of subjects into these two groups.

Consistent with our previous results in the DA treatment, if we only looked at the last 25 rounds, both groups of subjects exhibit very stable thresholds of similar magnitude, making them indistinguishable.

[22] The online appendix can be found at: https://econweb.ucsd.edu/~itrevino/pdfs/online_appendix_rt.pdf.

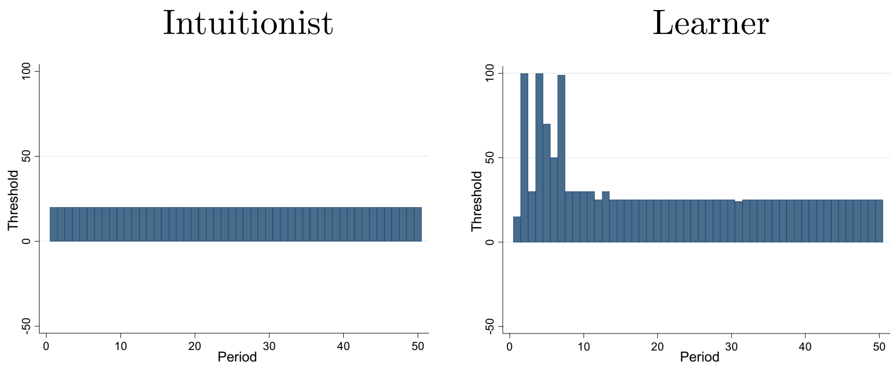

Fig. 11 Examples of evolution of thresholds in SM treatment, by group

5 Discussion: related models

In this section we discuss three existing models that could potentially make predictions for our data by relating high RTs to choices between two alternatives that have similar valuations to the decision maker. We investigate the predictions of these three models for our experiment and find some, but limited, support for them for the aggregate data (pooled across subjects), and mixed evidence for individual behavior, which is what we aim at understanding. More precisely, we will demonstrate that each of the models we describe below makes an identical qualitative prediction, which is that RTs should be decreasing in the distance between the signal that a subject receives and his threshold. Put differently, individual RTs should be a concave function of the signals received by a subject, with a maximum for signals that coincide with the observed threshold. It is this prediction at the individual level that fails in our data.

The first model is the Drift Diffusion Model (DDM; Ratcliff 1978; Ratcliff and McKoon 2008), which is a widely used model in psychology and neuroscience that studies the way in which the brain compares values to make binary choices. One of the key outputs of this model is the RT of subjects in these tasks. This model assumes that decisions are made by a noisy process that accumulates information over time from a starting point toward one of the two responses (or boundaries), and a response is chosen once one of these boundaries is reached (Ratcliff and McKoon 2008). The rate of accumulation of the information is assumed to be determined by the quality of information extracted by a stimulus. Ratcliff and McKoon (2008), for example, use a motion discrimination task where the stimulus is composed by a set of dots in a circle and, in each round, a proportion of the dots moves coherently either to the left or to the right, and the rest of the dots move in a random direction. The task for subjects is to decide in which direction the coherent dots move. When varying the proportion of dots that move coherently, they find that higher RTs are associated to higher levels of difficulty (i.e. low coherence) and to an almost equal probability of choosing the right and the wrong direction. On the other hand, when a high proportion of the dots move coherently, subjects make the right choice

Table 9 Summary statistics of differences in individual reported thresholds in the first 25 rounds in the SM treatment, by group

	Median	SD	N
Learners	6	16.36	7
Intuitionists	0	2.51	9

more often and they exhibit lower RTs. What this effectively means is that higher RTs arise as decisions become harder for subjects because they cannot clearly assess what is the right choice, given the information presented to them. So for Ratcliff and McKoon (2008) a stimulus that is more coherent is one that gives subjects a better idea of what choice to make.

Mapping the DDM to our experiment is not as straight forward a task as one might think, since there are some differences between our experiment and the typical DDM experiment. One clear difference is that the DDM studies individual decision making, while our experiment involves strategic interaction. Another important difference is that in a typical DDM experiment subjects are faced with an environment where each trial in the experiment is independent from the last, so there is no carry over between trials. This is true in the original Ratcliff (1978) experiments as well as more recent papers (e.g. Milosavljevic et al. 2010; Krajbich et al. 2010; Clithero 2018), where choices are made in a value-choice context. In our experiment, quite the opposite is true, since we present subjects with a learning task (they need to learn the best threshold to use), which involves arriving at the right set of expectations about the true state, but also about their opponent and his strategy. Information from previous trials is essential in this task and hence the trials are not independent from one another. A valid application of the DDM model to our context would therefore need to be a dynamic model where, based on previous experience, a subject updates the starting value of the DDM process. Such a model is beyond the scope of this paper. Finally, another important difference between our study and DDM studies is that we use RT to study individual learning processes, i.e. we study how each subject learns to set a threshold as he moves across rounds receiving different signals.

However, the DDM can, in a limited way, be used to think about the task facing subjects that we describe as Intuitionists. Consistent with our view of Intuitionists, the DDM can predict RTs in situations where, for example, subjects have to assess whether a certain number is higher than a fixed reference number (the Intuitionists' implicit threshold). According to the model for memory retrieval (Ratcliff 1978), subjects might have a hard time remembering a reference number, and if they receive a stimulus in the form of another number and have to decide whether this stimulus is higher than the reference number, the DDM predicts higher RTs for numbers that are closer to the reference number, since they require subjects to make a higher effort when assessing its value with respect to the reference (i.e. for very high or very low numbers it is "easier" to decide that they are higher or lower than the reference number). In our context, we think of Intuitionists as subjects who hold a threshold in their head which they try to retrieve once they are presented with a signal in a given round. Using the DDM language, a more coherent signal could be an extreme value, either very high or very low, and very far from the subject's threshold, which should imply an easy choice for subjects and hence a low RT. Likewise, a less coherent signal would be one for which subjects cannot easily assess which action to take (in our case, because it is close to their eventual threshold), and this would be associated to a higher RT.

Bearing this in mind, the DDM would predict that subjects in our experiment should exhibit longer RTs as they receive signals closer to their personal thresholds and that RTs should decrease as signals get further away from their personal thresholds. In this sense, we should expect the relationship between signals and RTs to be concave with a maximum at the observed threshold for each subject.

A similar prediction arises from the model of Gabaix and Laibson (2005) and Gabaix et al. (2006), that is tested in the experiment of Chabris et al. (2009). These papers propose an optimization theory called the Directed Cognition Model (DCM), which is based on dynamic programming and assumes that agents have limited cognitive resources. When time is a scarce resource, the DCM predicts that agents will allocate more decision time to choices between options of similar expected utility than to choices between options of dissimilar utilities. Just as the DDM, the DCM would predict, in the context of our paper, an inverse relationship between RT and the distance between signals and thresholds.

It is important to emphasize, however, that the DDM and DCM are designed to analyze individual decision problems, and not games.

The third model that we present here is an interpretation, in terms of RT, of the canonical global games model, as presented in Sect. 2. From Eq. 1 note that, by definition, when an agent observes a signal that has exactly the same value as the threshold, he does not have any strict preference over actions. This effectively means that for these signals agents are not sure about which action would yield a higher expected payoff. Interpreting this condition in terms of RT, a subject should exhibit a higher RT when confronted with signals that are closer to the subject's threshold than when observing signals that are far from it. This model, as the other two theories, would predict a concave RT function for each subject, with a maximum at the threshold chosen by them.

In summary, each of the models described above predicts that our RT data should exhibit a negative and significant relationship between RT and the distance between signals and future individual thresholds. We explore this prediction by performing a random effects OLS regression for the data, pooled across subjects, that has RT as the dependent variable and the difference (in absolute value) between the signal associated to each RT and the individual threshold that each subject converges to in the last 25 rounds as the independent variable.

Bearing our caveats in mind about the applicability of the DCM and DDM models to our data, we explore this prediction. In the DDM and the DCM models this stylized fact should describe behavior in all rounds of the experiment, while in our analysis, this effect should be stronger before the BPRT is reached. This is expected to be true whether the subject is classified as a Learner or an Intuitionist.

In order to investigate this, we run a series of OLS random effects regressions, reported in Table 10. In specification 1 we have the logarithm of RT as the dependent variable and the difference (in absolute value) between the signal and the threshold as the independent variable. We observe a negative and significant relationship between these two variables, just as hypothesized by the DDM and DCM. In specification 2 we add a variable for the round in which each decision takes place because we know that, on average, RTs decrease as we move across rounds, and find that the previously established relationship still holds, which again supports the DDM and DCM models. Specification 3 is similar to specification 2, with the addition of a coefficient for the Period squared, since in later rounds there is a flattening of RTs.

The results in specifications 1–3 support this stylized fact (negative and significant coefficient for the variable |signal − threshold|), and thus for the predictions of the DDM and DCM. Note that these regressions are run on data generated by all of our subjects, i.e., both Learners and Intuitionists. The fact that we get significant results, therefore, is notable since we do not necessarily posit that the relationship will hold for Learners. To separate our conjecture from the DDM and the DCM we run specifications 4–6, where we include two more independent variables. One is a dummy that takes the value of 1 for periods up to the BPRT (D(BPRT) in the table), and zero for the remaining periods, and the other variable is an interacted term that multiplies this dummy to the absolute value difference between the observed signal and the individual threshold. When we introduce this control for the periods before and after the BPRT, we find that the influence of the signal (or its distance from the threshold) is twice as large (on average) for the periods up to the one corresponding to the BPRT than after (comparing the coefficients for |signal − threshold|× D(BPRT) to |signal − threshold| in specifications 4–6), which supports our conjectured behavior. In other words, subjects seem to exhibit periods of stronger active consideration before their BPRT period than after, supporting our interpretation that the BPRT is a period of discovery. Therefore, we find only partial evidence for the DDM and DCM predictions, and we find that our RT observations are meaningful in terms of these predictions. An implication of this result is that it is possible that two different decision processes are used by a subject in the same experiment.

5.1 Strategic considerations

Studying RT in a game might bring a layer of complexity that is not present in individual decision making tasks. In our setup, for example, one might believe that subjects who fail to coordinate with their partners across the different rounds of the experiment are more likely to exhibit longer RTs on average. In other words, since subjects are not sure about the action that their partner will take, in every round it takes them a longer time to make a decision because they are constantly trying to adjust to their changing beliefs about their opponent's strategy.

To explore this possibility we look at the data at the pair level. We create a measure of convergence of behavior within a pair by taking the absolute value difference between the thresholds in the last 25 rounds of each pair member. We study whether a higher distance between the thresholds in a pair (less convergence) leads to higher

Table 10 RT as a function of the distance between signals observed and individual thresholds

	1	2	3	4	5	6
| signal - threshold |	– 0.013***	– 0.013***	– 0.014***	– 0.007***	– 0.006**	– 0.008***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
| signal - threshold |× D(BPRT)				– 0.014***	– 0.016***	– 0.014***
				(0.004)	(0.004)	(0.004)
D(BPRT)				2.635***	1.836***	1.66***
				(0.255)	(0.283)	(0.285)
Period		– 0.122***	– 0.352***		– 0.082***	– 0.278***
		(0.01)	(0.045)		(0.013)	(0.047)
Period^2			0.008***			0.007***
			(0.002)			(0.002)
Constant	4.63***	6.25***	7.438***	3.538***	4.953***	6.074***
	(0.17)	(0.218)	(0.317)	(0.203)	(0.306)	(0.401)

Clustered (by subject) standard errors in parentheses

*Significant at 10%; ** significant at 5%; *** significant at 1%

RTs. Table 2 in the online appendix reports the results of a similar regression to specification 2 above, but where we account for the distance between thresholds in a pair. We find no such effect.

We do a similar analysis looking at the types that we have described above, Intuitionists and Learners, since the behavior of the former is more stable and does not seem to be adjusted from round to round, while the behavior of the latter is the opposite. In that sense, it is important to understand if these types are endogenous in the sense that the matching protocol affects whether a subject becomes an Intuitionist or a Learner. In other words, is it the case that Intuitionists exhibit stable behavior from the get go because they are more likely to be matched with an Intuitionist and learn their partner's strategy early on, while Learners might take longer to set a threshold because they are matched with other Learners like themselves. We do not find any evidence that is consistent with this type of reasoning. In particular, over 76% of pairs are composed by one subject who is classified as an Intuitionist and one who is classified as a Learner, and the median difference in threshold behavior within pairs are not statistically different from those of the pairs that are constituted by subjects who are both Intuitionists or Learners. Likewise, the median difference in threshold behavior within pairs are not statistically different for pairs of Learners and pairs of Intuitionists.[23]

[23] This lack of statistical significance is expected since there are very few data points to have sufficient power.

6 Conclusion

In this paper we have attempted to gain insights into the thought process of subjects engaged in global games using the response times of their decisions. Quite remarkably, we have found that by looking at the highest or second highest response time exhibited by subjects in the early rounds of the experiment we can predict the eventual threshold they use in future rounds. This result is rather striking since response times are used not only as a way to gain qualitative insights into how different choices are represented in the decision making process, but rather as a tool to predict future choices. We know of few papers that attempt to do this.

We have shown that the accuracy of response time estimates as predictors of future behavior is comparable to that of choice data in early periods. Moreover, we have shown that these two measures are actually complementary since response time estimates provide additional predictive power over and above choice data.

In addition, we have presented evidence that these high response times represent different thought processes for different types of subjects. Based on the best predictor among the two highest response times, we classify subjects as Intuitionists and Learners and differentiate their behavior in the initial rounds of the experiment. This classification allows us to understand the different reasoning processes that lead different subjects to choose a similar threshold. That is, if one were to only look at choice data these two groups would be indistinguishable in terms of the thresholds they eventually set. In this sense, studying response times gives us an additional insight into the thought process that leads to setting a strategy in a global game.

We have also presented evidence in support of our distinction between Intuitionists and Learners in an out-of-treatment exercise. We observe behavior consistent with these two different groups in this new treatment, which illustrates that our categorization of subjects into types might be meaningful.

Finally, we look at the predictions of alternative models of cognition in the context of our paper and find mixed evidence about their predictions when analyzing behavior at the individual level.

In short, our paper provides an interesting insight into the usefulness of response times in the explanation of choice in global games.

Acknowledgements We would like to thank the National Science Foundation via Grant SES-105962 and the Center for Experimental Social Science at New York University for financial support. We would also like to thank Colin Camerer, Paul Glimcher, David Laibson, Ross Metusalem and Elizabeth Schotter for their advice.

References

Alos-Ferrer, C., Granic, D. G., Kern, J., & Wagner, A. K. (2016). Preference reversals: Time and again. *Journal of Risk and Uncertainty*, *52*(1), 65–97.

Busemeyer, J. R. (1985). Decision making under uncertainty: A comparison of simple scalability, fixed-sample, and sequential-sampling models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *11*(3), 538–564.

Caplin, A., & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, *105*(7), 2183–2203.

Carlsson, H., & van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, *61*(5), 989–1018.

Chabris, C. F., Laibson, D. I., Morris, C. L., Schuldt, J. P., & Taubinsky, D. (2009). The allocation of time in decision-making. *Journal of the European Economic Association*, *7*(2/3), 628–637.

Clithero, J. A. (2018). Improving out-of-sample predictions using response times and a model of the decision process. *Journal of Economic Behavior & Organization*, *148*, 344–375.

Fehr, E., & Rangel, A. (2011). Neuroeconomic foundations of economic choice—Recent advances. *Journal of Economic Perspectives*, *25*(4), 3–30.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Fudenberg, D., Strack, P., & Strzalecki, T. (2018). Speed and the optimal timing of choices. *American Economic Review*, *108*, 3651–3684.

Gabaix, X., & Laibson, D. (2005). *Bounded rationality and directed cognition*. Cambridge: Harvard University.

Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, *96*, 1043–1068.

Gul, F., & Pesendorfer, W. (2008). The case for mindless economics. In A. Caplin & A. Shotter (Eds.), *The foundations of positive and normative economics*. Oxford: Oxford University Press.

Heinemann, F., Nagel, R., & Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, *72*(5), 1583–1599.

Konovalov, A., & Krajbich, I. (2019). Revealed strength of preference: Inference from response times. *Judgment and Decision Making*, *14*(4), 381–394.

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292–1298.

Krajbich, I., Oud, B., & Fehr, E. (2014). Benefits of neuroeconomic modeling: New policy interventions and predictors of preference. *American Economic Review: Papers and Proceedings*, *104*(5), 501–506.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction times of value-based choice under high and low time pressure. *Judgment and Decision Making*, *5*, 437.

Morris, S., & Shin, H. (2003). "Global games: Theory and applications. In *Eight world congress: Advances in economics and econometrics: Theory and applications* Cambridge: Cambridge University Press.

Morris, S., & Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, *88*(3), 587–597.

Piovesan, M., & Wengstrom, E. (2009). Fast or fair? A study of response times. *Economic Letters*, *105*, 193–196.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.

Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal*, *117*, 1243–1259.

Rubinstein, A. (2008). Comments on neuroeconomics. *Economics and Philosophy*, *24*, 473–483.

Rubinstein, A. (2013). Response time and decision making: An experimental study. *Judgement and Decision Making*, *8*(5), 540–551.

Spiliopoulos, L., & Ortmann, A. (2014). The BCD of response time analysis in experimental economics. *Experimental Economics*, *21*, 383–433.

Szkup, M., & Trevino, I. (2019). Sentiments, strategic uncertainty, and information structures in coordination games, mimeo.

Webb, R. (2019). The (neural) dynamics of stochastic choice. *Management Science*, *64*(1), 230–255.

Wilcox, N. T. (1993). Lottery choice: Incentives, complexity, and decision time. *The Economic Journal*, *103*, 1397–1417.
Woodford, M. (2014). Stochastic choice: An optimizing neuroeconomic model. *American Economic Review: Papers and Proceedings*, *104*(5), 495–500.