Management Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Using Response Times to Infer Others' Private Information: An Application to Information Cascades

Cary Frydman, Ian Krajbich

informs
http://pubsonline.informs.org/journal/mnsc
MANAGEMENT SCIENCE
Vol. 68, No. 4, April 2022, pp. 2970–2986
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Using Response Times to Infer Others' Private Information: An Application to Information Cascades

Cary Frydman,[a] Ian Krajbich[b,c]

[a] Marshall School of Business, University of Southern California, Los Angeles, California 90089; [b] Department of Economics, The Ohio State University, Columbus, Ohio 43210; [c] Department of Psychology, The Ohio State University, Columbus, Ohio 43210

Contact: cfrydman@marshall.usc.edu, https://orcid.org/0000-0002-2122-376X (CF); krajbich.1@osu.edu, https://orcid.org/0000-0001-6618-5675 (IK)

Received: September 5, 2019
Revised: March 24, 2020; November 30, 2020
Accepted: December 6, 2020
Published Online in Articles in Advance: September 28, 2021

https://doi.org/10.1287/mnsc.2021.3994

Abstract. The standard assumption in social learning environments is that agents learn from others through choice outcomes. We argue that in many settings, agents can also infer information from others' response times (RT), which can increase efficiency. To investigate this, we conduct a standard information cascade experiment and find that RTs do contain information that is not revealed by choice outcomes alone. When RTs are observable, subjects extract this private information and are more likely to break from incorrect cascades. Our results suggest that in environments where RTs are publicly available, the information structure may be richer than previously thought.

History: Accepted by Yan Chen, decision analysis.

Funding: C. Frydman and I. Krajbich both acknowledge support from the National Science Foundation [Grants 1554837 and 1749824, respectively], and I. Krajbich acknowledges the Cattell Sabbatical Fund and the Ohio State University Decision Sciences Collaborative.

Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc.2021.3994.

Keywords: information cascade • herding • response time • reaction time • drift diffusion model • sequential sampling model • information aggregation • social learning

"Doubt is a creature within the air. It grows when someone hesitates." —Toba Beta

1. Introduction

In many economic environments, agents can acquire private information and also learn from the actions of others. For example, when making a recommendation to buy a company's stock, an analyst can draw on her own private research while also relying on the public recommendations of other analysts (Welch 2000). The seminal theoretical works of Banerjee (1992) and Bikhchandani et al. (1992) show that after the consensus of public recommendations is strong enough, agents will optimally ignore their private information and "follow the herd." This failure of private information aggregation is caused by the limited set of public actions from which agents can learn (e.g., a "buy" or "sell" recommendation but nothing in between).

Lee (1993) shows that full information aggregation can be restored if the action set is continuous—and is therefore rich enough to transmit private signals to others. We argue that the action set is enriched when features of the choice *process* are observable, such as the amount of time it takes to make a decision. For example, consider a labor market where a job applicant receives a job offer and observes how long it took to receive the offer after the interview. The speed with which an offer is extended can carry information about the strength of the employer's preference for the applicant, and this extra information can influence the applicant's decision to accept the offer. Indeed, Becker et al. (2010) find that job applicants are more likely to accept a job offer if it is delivered quickly (controlling for *ex post* employee performance ratings). The intuition is that a fast offer signals that the firm has a strong preference for the candidate. Van de Calseyde et al. (2014) provide similar evidence in a laboratory setting where decision times are exogenously manipulated.[1] From a theoretical perspective, if response times (RTs) reveal an agent's private information, then many economic environments will contain richer information structures than are typically modeled.

In this paper, we conduct a carefully controlled laboratory experiment to test whether the provision of RT information has an impact on behavior during social learning. In order for RTs to have an impact on behavior, two conditions must be satisfied. First, RTs must at least partially reveal an agent's private information. Second, agents must be able to correctly invert the mapping from private information to RT. To guide our study of this relationship, we rely on models of optimal time allocation from mathematical psychology and more recently, economics (Woodford 2014, Fudenberg et al. 2018). These models are called sequential sampling models (SSMs; of which the drift-diffusion model (DDM) is a prominent example), and we review them in Section 2.

Our experimental setting is based on the classic design of Anderson and Holt (1997). There is an uncertain binary state of the world, and each subject receives a private and conditionally independent signal. Subjects move sequentially and provide a prediction of the state of the world, which then becomes public information to all subjects (see Weizsäcker 2010 for a meta-analysis of this class of experiments). Therefore, before a subject makes his prediction, he has access to his private signal and all actions of the subjects that have moved before him—but not their private signals. If these binary actions are the only observables, the theoretical prediction is that information aggregation will quickly cease, and with positive (but often small) probability, agents will herd on the wrong action. We study two versions of this paradigm, one where only the choice outcomes are observable and one where RTs are also observable.

The literature on sequential sampling models indicates that longer RTs are typically associated with more conflicted decisions (see Section 2). We therefore expect that when a subject's private information conflicts with the actions of his predecessors, he should take more time to decide. Crucially, if the subsequent subject in line is aware of this mapping between conflict and RTs, then he should infer that his predecessor was more conflicted if the choice was associated with a longer RT. Interestingly, a footnote in the original Anderson and Holt (1997) paper hints at the idea that decision *process* data can reveal private information: "In an admittedly uncontrolled demonstration experiment … students seemed to use visual and voice cues in an attempt to discern whether the person making a prediction was agonizing over a sample draw that seemed unlikely given the pattern of earlier public predictions" (Anderson and Holt 1997, p. 851). By experimentally varying the ability of subjects to observe the RTs of their predecessors, here we test whether behavior is causally affected by the decision process information of others.

The two main results of the paper can be summarized as follows. First, we find that RTs carry information about both the public information and the subject's private signal: longer RTs are associated with decisions in which a subject's private signal conflicts with the actions of the group. Importantly, RTs contain additional information that is *not* contained in the choice outcome alone. Subjects are indifferent between choosing with their private signal when they have observed a net imbalance of two previous actions against their own private signal (in line with previous cascade experiments (Weizsäcker 2010)). On these "indifference" trials, subjects exhibit their longest RTs. As we move away from indifference, RTs decline monotonically, consistent with these decisions becoming easier. An action that is consistent with the private signal and two predecessors' actions is faster than an action consistent with the private signal and one predecessor's action, which in turn is faster than an action that is only based on the private signal. This result is important because for a substantial number of trials (45%), choice probabilities do *not* vary as we move away from the indifference point. Thus, RTs reflect the public information, even after controlling for choice outcomes. More importantly, for a given sequence of predecessor actions, subjects who follow the majority do so faster when their private signal aligns with the majority compared with when it does not. Thus, RTs also reflect private information. In this basic sense, variation in continuous RTs reflects information that is obscured in the discrete choice data.

Our second main result is that subjects are able to extract information about predecessors' private signals from their RTs. This indicates that subjects are able to invert the mapping from private information to RTs. Specifically, when a subject's private signal conflicts with his predecessors' moves, the subject is more likely to choose with his own private signal if the predecessor exhibited a long RT. That is, a subject is more likely to go against the choice outcome generated by a slow decision from his predecessor, compared with a fast decision.

Taken together, our two main results provide support for the joint hypothesis that (i) longer RTs are associated with more conflicted decisions and (ii) subjects are aware of this mapping. We also find that RTs are very similar across our two experimental conditions, suggesting that subjects are not strategically manipulating their RTs or engaging in "cheap talk."

A natural question is whether the provision of RTs also generates an increase in efficient outcomes. After all, given that subjects are able to infer additional information from their predecessors, this should generate higher payoffs when RT information is available. We do not find a significant increase in payoffs when providing subjects with RT information. However, this result is not surprising when looking at the data unconditionally. To see why, note that an increase in efficiency is most likely to occur during incorrect cascades (those in which the first two players receive incorrect signals). This is because RT information is most valuable when it can be used to stop an incorrect cascade, and by design, such incorrect cascades occur with low probability.[2] If we restrict our analysis to incorrect cascades, we do find that the provision of RTs generates a modest increase in payoffs.

Although we find that RT information does systematically impact choices, many of our basic choice results are consistent with the past experimental literature on information cascades. In particular, we find systematic deviations from Bayesian Nash equilibrium (BNE) as subjects in our experiment continue to condition their behavior on a predecessor's move even during an information cascade (Anderson and Holt 1997; Kübler and Weizsäcker 2004, 2005; Goeree et al. 2007; Weizsäcker 2010).[3] Many theories have been proposed to explain such deviations including quantal response equilibrium (McKelvey and Palfrey 1995, Goeree et al. 2007), level k (Stahl and Wilson 1994, Crawford and Iriberri 2007), cognitive hierarchy (Camerer et al. 2004, Coricelli and Nagel 2009), and more recently, naïve herding. The naïve herding theory argues that subjects take their predecessors' actions at face value and do not infer that such information may be redundant given the actions of earlier predecessors (Eyster and Rabin 2010, 2014). A recent experimental paper provides support for this theory, demonstrating that deviations from Bayesian Nash behavior in social learning settings are likely to stem from strategic naiveté rather than errors in Bayesian updating (Eyster et al. 2015). Our paper examines distinct questions about whether non-choice data can reveal private information and how subjects condition on this non-choice data when making decisions in a strategic setting.

Finally, it is important to emphasize that the RT we analyze in this paper is a measure of the amount of time that it takes for a subject to execute a decision. Specifically, we define RT as the time that elapses from the experimenter's revelation of a subject's private signal until the subject reveals his decision. This is fundamentally different from the time that elapses *between* decisions made by different agents, which is studied in social learning models where decision opportunities arrive at random times (Guarino et al. 2011, Herrera and Hörner 2013). In such settings, agents can use the arrival rate of past decisions to infer information about the state of the world—without conditioning on the RT of each individual agent. Our approach is also different from the literature on strategic delay (Chamley and Gale 1994, Gul and Lundholm 1995, Zhang 1997) because in our setting, agents cannot choose the order in which they make their decisions. Furthermore, as we describe in Section 3, an agent's RT can be viewed as the optimal solution to a statistical decision problem, and thus, it cannot be strategically manipulated without cost.

2. Background: A Framework for Studying Response Times

Our hypothesis that RTs can reveal private information in a strategic setting is motivated by a rapidly growing literature on the economic analysis of RTs.[4] There are two separate strands of this literature. One strand has argued that long response times are generated by employing high effort to make an accurate decision (Wilcox 1993, Eliaz and Rubinstein 2014, Rubinstein 2016). The second strand has argued that long RTs are generated by difficult decisions in the sense that the options under consideration have similar utility (Busemeyer and Townsend 1993, Krajbich et al. 2010, Mormann et al. 2010, Philiastides and Ratcliff 2013, Polanía et al. 2014, Rodriguez et al. 2014).

Interestingly, both of these phenomena can be produced by an optimal decision rule for a Bayesian agent collecting stochastic information to guide his choice (Fudenberg et al. 2018). To see this, it is useful to review a framework for analyzing the joint distribution of choice outcomes and response times. This framework has been extensively used in the mathematical and cognitive psychology literatures for decades, and the resulting models are known as SSMs.

SSMs describe the dynamics of the decision process, from the time that a subject is presented with two alternatives until the time he executes a decision. For example, suppose a subject is asked to choose between option A and option B, which represent the optimal actions in two different states of the world. In an SSM, there is a latent decision value Z that is computed by stochastically sampling over time (within a trial) the difference in subjective values between option A and option B ($u(A)$ and $u(B)$, respectively). This variable Z captures the accumulated "net evidence" favoring one option over the other, and its evolution can be described with the following equation:

$$Z_t = \mu t + B_t, \tag{1}$$

where B_t is a standard Brownian motion variable and $\mu \equiv (u(A) – u(B))$ is the rate of evidence accumulation, known as the drift rate. In our particular setting, the drift rate is a function of both public and private information. This process of evidence accumulation continues until Z_t crosses a boundary, $b_t \equiv \pm\delta$, triggering an action, $a \in \{A, B\}$. Specifically, the agent chooses A if $Z_\tau \geq \delta$ and B if $Z_\tau \leq -\delta$, where the RT is $\tau = \inf\{t : |Z_t| \geq b_t\}$.

Here, we have assumed that the boundaries for actions A and B are symmetric and constant over time, as in the simplest DDM. In other applications, the boundaries can be asymmetric to encode prior information favoring one action over the other or can collapse over time to reflect the urgency of a decision (Fudenberg et al. 2018).

The most "difficult" decisions are those where $\mu = 0$ (i.e., when there is no drift toward either boundary and the sole determinant of RT is the Brownian noise). One can then use the well-known result that the mean RT, conditional on the drift rate, boundary, and noise parameter, is given by $E[RT] = \alpha tanh(\alpha\beta)$, where $\alpha = \frac{\delta}{\mu}$, $\beta = \left(\frac{\mu}{\sigma}\right)^2$, and σ is the standard deviation of the Brownian noise. This implies that the mean RT is maximized at $\mu = 0$. Intuitively, the average RT is longest when "evidence" accumulation is driven only by noise.

Note that because of the intrinsic noise in the process, there is a positive probability that the option with lower subjective value is chosen. These "mistakes" can be minimized by choosing the boundary values $\{\delta\}$ to be sufficiently large, but this comes at the cost of longer RTs. The boundary values therefore parameterize the so-called "speed-accuracy trade-off". Under certain assumptions this algorithm implements the optimal solution to the speed-accuracy problem, in the sense that it maximizes the probability of choosing the higher-valued option for a desired RT (Wald 1945, Bogacz et al. 2006, Krajbich et al. 2014, Woodford 2014, Fudenberg et al. 2018).

With this framework in hand, we can now describe how it generates both fast and slow mistakes, depending on the model parameters. Consider the first claim that longer RTs result in more accurate decisions. For a given level of decision difficulty, $|u(A) - u(B)|$, a subject can increase his expected accuracy by increasing the distance between the two boundaries. This will however come at the cost of longer RTs because it will take the process a longer time to reach one of the two boundaries. Hence, for a given level of decision difficulty, more accurate decisions are associated with longer RTs.

Now take the second claim, which is that longer RTs are associated with more difficult decisions. Because decision difficulty decreases in $|u(A) - u(B)|$, then for a given set of boundaries $\{-\delta, \delta\}$, more difficult decisions will take longer to reach one of the two boundaries. Hence, more difficult decisions are associated with longer RTs. As difficulty increases (as $|u(A) - u(B)| \to 0$) and the agent approaches indifference, he chooses randomly but counterintuitively, takes the most time to do so (Mosteller and Nogee 1951, Moffatt 2005, Chabris et al. 2009, Dickhaut et al. 2009, Krajbich et al. 2014, Woodford 2014, Konovalov and Krajbich 2017, Fudenberg et al. 2018).

In many environments, including the social learning experiment studied in this paper, it is of course possible that both decision difficulty and the boundary magnitude δ vary across trials and agents. It therefore becomes an empirical question to assess which force dominates in generating the RT distributions. Perhaps more importantly for this paper, it is also an empirical question whether subjects are aware of which force is dominant.

Finally, it is important to note that one implication of the SSM's speed-accuracy trade-off is that RTs cannot be treated as cheap talk (i.e., they cannot be arbitrarily manipulated without a cost). Recall that the SSM is a solution to a statistical decision problem that minimizes mean RT for a given accuracy level. Hence, a subject cannot strategically decrease his RT without sacrificing expected accuracy. Moreover, our experimental design removes any incentive for subjects to engage in cheap talk because payoffs are independent of the actions of other players.

3. Experimental Design

Our experimental design is inspired by the canonical information cascade experiment in Anderson and Holt (1997). The main unit of analysis is defined as a "round," where a group of eight subjects was arranged in random order and incentivized to predict the uncertain state of the world. The order of subjects is indexed by $i = 1, 2, \ldots 8$. In each round, there were two possible states of the world, $S \in \{A, B\}$, and the prior probability of each state was 0.5. The subject's task was to determine which state of the world was more likely. On his turn, subject i received a conditionally independent private signal $s \in \{a, b\}$ such that $\Pr(s = a|S = A) = \Pr(s = b|S = B) = 2/3$, and he also saw the actions of all previous subjects $1, 2, \ldots i-1$. The subject then had 10 seconds to predict the correct state of the world. Following the incentive structure of earlier experiments on information cascades, the subject was paid \$1 if his prediction was correct and \$0 otherwise.

At the beginning of each round, the first subject saw his private signal and had 10 seconds to make a choice. At the end of the 10 seconds, that subject's choice was displayed to every member of the group. Subjects observed this new public information for as long as they liked and then clicked on a "continue" button. After every subject pressed the "continue" button, the second group member received his private signal, and in addition, the first subject's choice was displayed again. As before, the second subject had 10 seconds to decide between A and B, after which his decision was revealed on a feedback screen to every group member. Each subsequent turn proceeded in this manner with subjects always receiving a table containing the previous choices from the current round. After all eight subjects made their choices, the true state was revealed, and those who predicted correctly received \$1 (and \$0 otherwise).[5]

All subjects played the game in two different conditions. In the "no-RT condition," the game was implemented exactly as described. In the "RT condition," the game was identical to the "no-RT condition," except the feedback table additionally contained each previous subject's RT. Thus, subjects in the RT treatment could also observe how long their predecessors took to decide. It is important to note that in both treatments, the decision stage always took exactly 10 seconds, regardless of the subject's RT.[6] This was

Figure 1. (Color online) Example Screenshots from the Experiment in the No-RT Condition (Upper Panel) and RT Condition (Lower Panel)

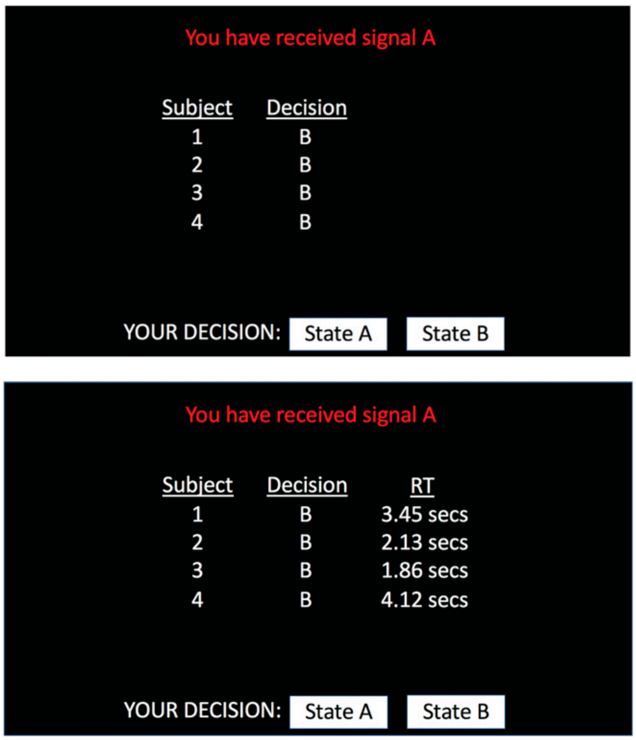

Notes. In this example, the subject is in the fifth position and readily has access to the choices (and RTs in the RT condition) of all four predecessors. The relatively slow decision of subject 4 might signal that he received an A signal. For illustration purposes, the size and location of the text and buttons have been slightly altered.

done so that subjects could only observe RT information in the RT condition while also equating the overall time spent in each condition. Figure 1 provides an illustration of the experimental display shown to subjects in each of the two conditions.

We chose to provide explicit measures of RT because our goal was to provide a sharp test of whether subjects use RTs in a manner that is consistent with SSMs. By explicitly providing subjects with RT information, we reduce the concern that subjects would be unable to measure their predecessor's RT; this enables a clean test of our hypotheses about how subjects respond to long versus short RTs in different situations. The explicit display of RT also mimics the way in which timing information is presented in many online marketplaces.

Seventy-two undergraduates from The Ohio State University participated in the experiment. The experiment was conducted in four sessions, the first three with 16 subjects each and the last with 24 subjects. We randomized the order of each condition across sessions: in the first and third sessions, subjects participated in 15 rounds of the no-RT condition followed by 15 rounds of the RT condition; this order was reversed in the second and fourth sessions. In each round, subjects were randomly rematched into groups of eight. By counterbalancing the order of conditions across sessions, we were able to additionally conduct between-subjects tests by analyzing only the first condition from each session. As we explain later in Section 5.4, this design feature is useful as it can help to reduce concerns about experimenter demand effects, which are most severe for within-subject tests.

At the beginning of each session, subjects gave informed written consent and then went through paper instructions for the relevant condition, including a short comprehension quiz. Experimenters verified the quiz answers and answered any questions before proceeding with the experiment. Subjects first completed two practice rounds without pay. During these practice rounds, the decision screen lasted 15 seconds. After completing the first condition (i.e., the first 15 paid rounds), subjects received a new set of paper instructions for the second condition. Again, the experimenters answered any questions before proceeding with the experiment. There were no quizzes or practice rounds for the second condition because the only difference was the addition/subtraction of the RT information. Each session lasted approximately two hours. The average payoff was \$20.14 (range: \$15–\$25, standard deviation: \$1.95) in addition to the \$5 show-up fee. The experiment was programmed in Multistage. The full experimental instructions are available in the online appendix.

4. Theoretical Predictions

4.1. Bayesian Nash Equilibrium

As a rational benchmark, we first describe the solution to the game using the BNE concept. This analysis follows directly from Anderson and Holt (1997). If the first player receives an a signal, then he will choose A, and if he receives a b signal, he will choose B. Thus, the first player's action reveals his private signal. If the second player's signal matches the first player's move, then the second player chooses with his signal. Alternatively, if the second player observes a signal that does not match his predecessor's move, then he infers that one a signal and one b signal have been drawn in total, and thus, his posterior that the state of the world is A is 1/2. In this case, we assume that the second player chooses with his own signal (such a tie-breaking assumption can be motivated by a positive probability that the first player makes an error).

If the third player observes both the first two players choosing the same action A, then he infers that both players received an a signal. In the case that the third player receives an a signal, then he also clearly chooses A. In the case that he observes a b signal, then he perceives a total of two *a* signals and one *b* signal. Thus, the third player chooses A, despite his private information that favors action B. In other words, the third player chooses *A* regardless of his private signal. This triggers an information cascade, with players i = 4, … 8 adopting the same logic and selecting action A regardless of their own private signal.

Taken together, the prediction of the BNE theory is that as soon as one state of the world receives on net two decisions in its favor, all subjects should choose this state of the world, regardless of their private signal. Importantly, this means that it is optimal for all participants to cease their belief updating at this point.

4.2. Deviation from Bayesian Nash Equilibrium

Although BNE theory makes sharp predictions in our setting, previous experiments using a similar design have documented systematic deviations from BNE predictions. To explain such deviations, many papers invoke behavioral theories of decision making in strategic settings, including quantal resonse equilibrium (McKelvey and Palfrey 1995), level k (Stahl and Wilson 1994), cognitive hierarchy (Camerer et al. 2004), and most recently, naïve herding (Eyster and Rabin 2010). Although we do not take a stand on the specific behavioral theory that generates deviations from BNE in the prior literature, it is still useful to describe a common prediction made by these alternative theories, namely that subjects continue to infer information from choices that are uninformative in the BNE framework. In particular, choices made during a cascade will continue to convey information and influence later subjects' behavior.

To operationalize deviations from BNE, we create a variable called the net public information (NPI), which is defined for subject i in each round as $NPI_i = (-1)^{1+1_{a,i}} \times \Sigma_{n=1}^{i-1}(1_{A,n} - 1_{B,n})$. The first term takes on the value of 1 if the subject receives an *a* private signal and −1 if he receives a *b* private signal. The second term gives the difference in the observed "A" moves and observed "B" moves. Under BNE, after $|NPI_i| > 1$, an information cascade is triggered, and the actions of subjects $i + 1$, $i + 2$, … 8 will not reveal their private information. In contrast, under several of the behavioral theories mentioned, a subject's action will depend on the value of NPI, even when $|NPI_i| > 1$. To guide our analyses of deviations from BNE, we use a simple rule that maps NPI to actions, where a subject's belief that his signal matches the state of the world is strictly increasing in NPI_i. Such a rule is consistent with, for example, an extreme version of the naïve herding model of Eyster and Rabin (2010). It is important to emphasize that we use NPI as a simple way to demonstrate the link between RTs, private signals, and choices. The specific construction of NPI will not affect our main empirical tests of whether RTs convey information about private signals and whether subjects use this information in their decisions.

In order to transform NPI into a choice probability, we assume a standard logit choice function. This assumption can be justified on the grounds that the DDM (a canonical SSM) implies that choices will follow a logistic distribution (Webb 2018) and that the probability of choosing in line with one's signal is strictly increasing in NPI. In particular, we assume that

$$P_i(choose\ \ signal) = \frac{1}{1 + e^{-\lambda \cdot (NPI_i - k)}}, \tag{2}$$

where λ parameterizes the noise variance in subjects' choices and k parameterizes the value of the *NPI* where a subject is indifferent between choosing with or against his signal.

5. Results

5.1. Inferring Information from Response Times

Because our main goal is to test whether subjects can infer information from the RTs of their predecessors, we first assess whether RTs do in fact contain information about private signals. To do so, we investigate whether subjects' choices and RTs were associated with the *congruency* between their private signals and others' choices, as captured by our NPI measure. We define a congruent signal as one in which NPI ≥ 0, and an incongruent signal as one in which NPI < 0. In other words, a private signal is congruent when it matches the majority of previous choices. We find that subjects were more likely to choose in line with their signal as the NPI increased (Figure 2). For positive values of NPI, subjects almost always chose with their signal (98%). Subjects were roughly indifferent (51%) about choosing with their private signal when NPI = –2. A logistic regression confirms the positive effect of NPI on choosing in line with the private signal (p < 0.001) and indicates an indifference point of NPI = −2.25.[7] These results are inconsistent with BNE theory, which predicts that a cascade should be triggered when | NPI | > 1. Under BNE, when NPI = –2, subjects should have always chosen against their signal. Instead, they only did so on 49% of trials, and as the NPI became more negative, subjects more frequently contradicted their private signal.[8]

Turning to the RT data, we find strong support for the prediction that RTs peak near indifference (Figure 3). Within the no-RT condition, the highest median RT occurred for an NPI of –3, with –2 being a close second (if we pool both no-RT and RT conditions, the highest median RT occurs at –2).[9] When conditioning on positive values of NPI, we find that there is no significant relationship between NPI and the probability that a subject chose with his signal (p = 0.34), and yet, RTs significantly decreased by 65 milliseconds (ms) for every unit of NPI (p = 0.056, using log(RT) as the

Figure 2. (Color online) Probability of Choosing in Line with One's Private Signal as a Function of the NPI (i.e., the Net Number of Prior Choices Aligned with the Private Signal)

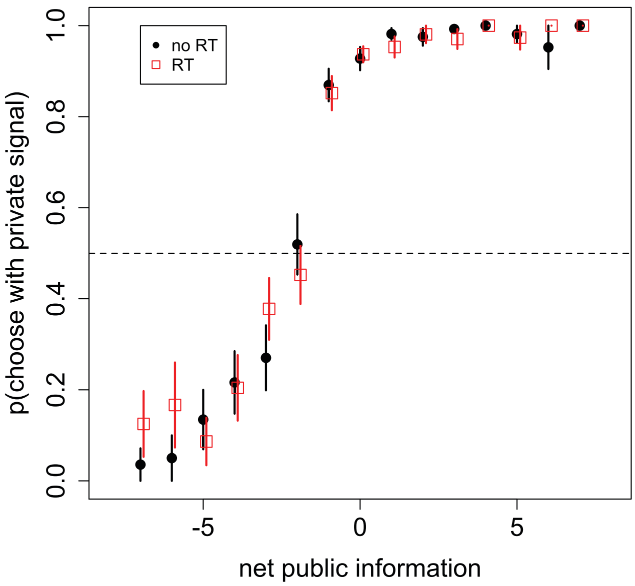

Notes. Dots represent data from the no-RT condition, and squares represent data from the RT condition. In both conditions, subjects are roughly indifferent between choosing with/against their private signal when there are net two prior choices inconsistent with the private signal. Vertical lines are standard errors of the mean, clustered by subject.

dependent variable in a linear regression on NPI). This

Figure 3. (Color online) Response Time as a Function of the NPI

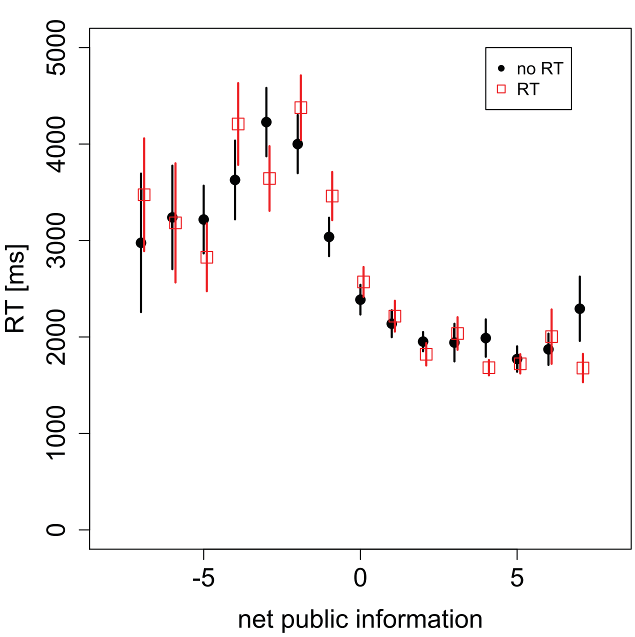

Notes. Dots represent data from the no-RT condition, and squares represent data from the RT condition. In both conditions, RTs peak at or near the indifference point of NPI = –2. Vertical lines are standard errors of the mean, clustered by subject.

Figure 4. The Average log(*RT*) of Subject *i* as a Function of the Net Number of Group Choices That Are Aligned with Subject *i*'s Choice

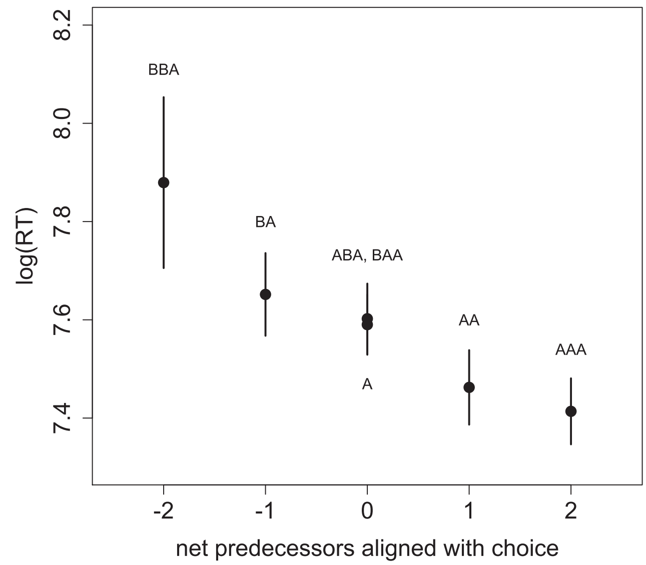

Notes. Data are limited to no-RT condition cases where subject *i* was in the first three positions of the sequence and chose in line with his signal. The labels in the plot indicate the particular scenarios in each data point, with position increasing from left to right (e.g., BA indicates a B choice for the first subject and an A signal/choice for the second subject) and with each trial coded so that A is the signal/choice of subject *i*. Vertical lines are standard errors of the mean, clustered by subject.

result provides our first piece of evidence that RTs contain information that is obscured in the choice data.

The previous analysis includes trials from all eight positions, and thus, we are pooling across situations with and without cascades (recall that a cascade cannot be triggered until the first two subjects have made their decisions). We therefore conduct an additional analysis restricted to trials in the first three positions; in these trials, a subject can be sure that his immediate predecessor chose in line with her private signal. We find a clear log-linear relationship between RT and the net number of prior choices aligned with the subject's signal (restricted to subjects who chose with their private signal) (Figure 4). This indicates that RTs reflect *public* information, conditional on following one's signal.

To investigate whether RTs specifically reflect *private* information, we analyze two separate cases, in trials where a cascade was potentially triggered (positions 4–8). First, consider the case where a subject chooses an action that is *against* the herd, which should occur only when his private signal is incongruent. In such circumstances, the choice outcome reveals the private signal, and RT should not contain any additional information about the (incongruent) private signal. Indeed, we find that on 92% of trials in which a subject chose against the herd, his choice matched his private signal.

Now consider the more interesting case where the subject chooses in line with the herd. In this case, there is more uncertainty about whether the subject received a congruent signal (and thus, had a trivial decision) or an incongruent signal (and chose against his private signal). On those trials where a subject chose in line with the herd, we find that the choice matched the private signal in only 66% of trials. Because choices contain less information about private signals when subjects choose in line with the herd, these are the trials where we would expect RTs to contain additional information about private signals.

To test whether RTs contain additional information beyond choice, we analyze trials in the no-RT condition where subjects chose in line with the majority of their predecessors' choices (87.5% of all trials with $|\text{NPI}| > 1$; 46.7% of all trials). We find that on trials where subjects received a congruent signal and thus, had an easy decision, the mean (median) RT was 1,914 (1,792) ms. In contrast, when a subject received an incongruent signal (and chose to follow the herd), the mean (median) RT was over twice as long: 3,869 (3,733) ms (paired t tests: $t(60) = 9.06$, $p = 10^{-12}$, $t(60) = 8.58$, $p = 10^{-11}$ for mean and median, respectively; 12 subjects who did not have observations in both cases were omitted from these t tests but not from the means/medians). This demonstrates that RTs do carry information about private signals, beyond the information contained in choice outcomes.

One potential concern with the analysis is that the NPI does not capture all publicly available information contained in the sequence of choice outcomes. In other words, two trials can be characterized by the same NPI, but the sequence of decisions that generate this NPI can differ. It is therefore important to consider the sequence of choices before we can conclude that RTs truly carry information that cannot be inferred from choice outcomes. To do so, we can further restrict our analysis to trials where *all* predecessors chose the same option (either A or B). Among this subset of trials, a sufficient statistic for the information contained in choice outcomes is given by the subject's position in the sequence (or equivalently, the length of the cascade). Thus, for a given position in the sequence, the only dimension on which RTs can systematically differ is the congruency of that subject's private signal.

Because SSMs predict that easier decisions are associated with shorter RTs, we expect that on trials where a subject receives an incongruent signal and chooses with the herd, RTs will decrease with the size of the herd. In contrast, on congruent signal trials where the subject chooses with the herd, there is no conflict between the private signal and the herd, and so, we expect the RT to be less affected by the size of the herd.

Consistent with both of these hypotheses, Figure 5 displays the median RTs on trials where subjects chose with the herd, conditional on every subject choosing the same option up to that point. The figure shows that the difference in median RT between incongruent and congruent signals was significant for all but the largest herd size of seven ($p = 0.01$, $p = 0.0009$, $p = 0.03$, $p = 0.08$, $p = 0.02$, and $p = 0.56$ for herd size = 2–7, respectively, based on two-sided t tests treating each subject as one observation). In other words, given the same history of choice outcomes and conditional on choosing with the herd, RTs are longer when private information conflicts with all previous choices.

Moreover, Figure 5 shows that on trials with a congruent signal (i.e., the subject chose with his private signal), RTs were independent of the size of the herd; however, on trials with an incongruent signal (i.e., the subject chose against his private signal), RTs were significantly longer but decreased with the size of the herd. A linear regression of log(RT) on size of the herd confirms these results: when regressing log(RT) on herd size for congruent signals, there was no significant effect ($p = 0.48$). However, there was a significant effect of herd size on log(RT) for incongruent trials ($p = 10^{-7}$) and also, a significant interaction effect between incongruent trials and herd size ($p = 0.006$), indicating that

Figure 5. (Color online) Response Time as a Function of Cascade Length, Conditional on All Predecessors Choosing the Same Option (No-RT Condition Only)

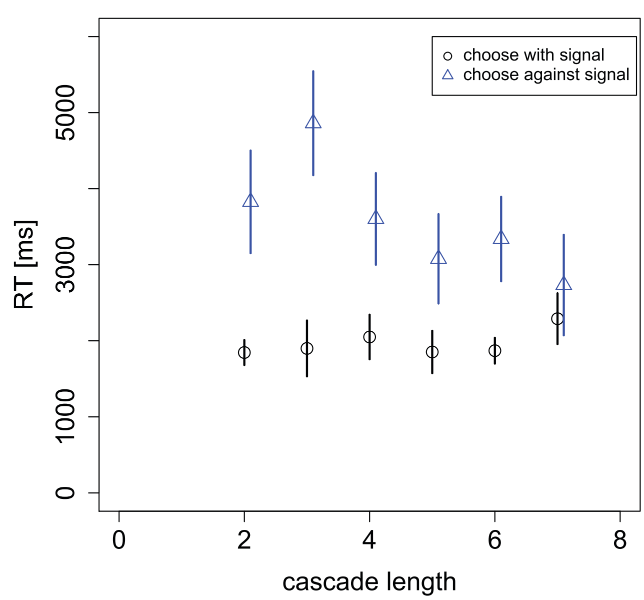

Notes. In these trials, the length of the cascade is equal to the subject's position in the sequence minus one. Circles represent cases where subjects chose in line with their private signals (i.e., they received a congruent private signal), whereas triangles represent cases where subjects chose against their private signals (i.e., they received an incongruent private signal). Vertical lines are standard errors of the mean, clustered by subject.

RTs are more sensitive to herd size in incongruent trials compared with congruent trials.

In summary, these analyses demonstrate that RTs provide additional information about subjects' private signals, beyond what is contained in their publicly observable choice outcomes. Before we test whether subjects can infer information from these RTs, it is first important to check whether subjects might have strategically manipulated their RTs in the RT condition. This is a critical check because strategic manipulation could alter the interpretation of the relationship between choices, RTs, and NPI.

To test this, we ran a Kolmogorov–Smirnov test on the difference in RT distributions between the RT and no-RT conditions. Specifically, we tested whether the two distributions displayed in Figure 3 are statistically different from one another. We found no significant difference between the two RT distributions overall ($D = 0.039$, $p = 0.39$), nor at the indifference point of NPI = –2 ($D = 0.112$, $p = 0.74$).

5.2. The Impact of Predecessors' Response Times on Behavior

How can a subject make use of the information contained in RTs, in order to maximize his payoff? In those situations where a subject's signal conflicts with his predecessors' choices, the subject should more heavily discount his predecessors' choices when they are slower. Intuitively, when a predecessor's choice is slow, this suggests that the predecessor received an incongruent signal.

Thus, one plausible strategy is as follows: when a subject receives a congruent signal, he always chooses with his signal. However, when a subject receives an incongruent signal, the probability that he chooses with his signal decreases with the speed of his predecessors' decisions. To test whether subjects conditioned their decisions on predecessors' RTs, we therefore analyze behavior on the subset of trials in which subjects received incongruent signals. We also restrict analysis to those trials after the third position because RT information was not pivotal for any subject in the first three positions.

We estimate a logistic regression where the dependent variable is *choose with private signal* (Table 1). The key variable is the triple interaction between the immediate predecessor's RT, the RT condition, and a dummy for whether the private signal conflicted with the immediate predecessor's move. The coefficient on this variable provides the (linear) effect that RT had on choosing with the private signal, in situations where the private signal conflicted with the predecessor's move (in the RT condition compared with the no-RT condition). In column (1), we find that this coefficient is significantly positive, indicating that when a subject's private signal conflicted with his predecessor's move, he was more likely to choose with his private signal as the predecessor's RT increased. In column (2), we find that the result is robust to the order in which the conditions were run.

One assumption in the regression model in Table 1 is that RT has a linear effect on behavior. To relax this assumption, we run a complementary test in which we assess whether subjects were more likely to choose with their signal when their predecessor made a "slow" versus "fast" decision. Here, we define a decision as "slow" if and only if its RT was longer than (or equal to) the median RT of 2,028 ms.[10] Figure 6 shows that in the RT condition, subjects contradicted their signal in 65.5% of trials when the predecessor's RT was slow, compared with only 33.1% of trials when the predecessor's RT was fast ($p < 0.001$). The figure also shows that there was a much smaller difference between these two probabilities in the no-RT condition (42.7% for slow decisions versus 38.0% for fast decisions, $p = 0.41$). The difference in fast versus slow decisions across conditions is highly significant ($p < 0.001$), which is consistent with regression results in Table 1. This indicates that in trials where a subject's private information conflicted with his predecessor's choice, he was twice as likely to contradict his signal when his predecessor made a fast decision. In Figure A.2 in the online appendix, we see similar results looking at earlier predecessors, who chose up to five positions before the subject.

5.3. Efficiency Differences

In this section, we investigate whether the availability of RT information translates into an increase in efficiency. To begin this analysis, we compute a theoretical benchmark for the increase in efficiency driven by RT information. To compute an upper bound on this increase, we assume that RTs perfectly reveal each subject's private signal. This implies that, when RT information is available, each subject's information set contains all predecessors' private signals. We ran a simulation to compare the efficiency in this case (under BNE) with the case where only the predecessors' actions are observable (the standard information cascade environment, which corresponds to our no-RT condition). Figure 7(a) plots the average payoff as a function of signal informativeness for those trials in which players receive an incongruent signal and are located after the third position in the sequence.

As expected, the provision of RT information leads to an increase in efficiency for all levels of signal informativeness. For our experimental parameter value of Pr(*correct signal*) = 2/3, the average payoff increases from approximately \$0.64 to \$0.68, which is a relatively small improvement given full revelation of all private signals. Consistent with this small predicted improvement, in the same subset of our experimental data we

Table 1. Dependent Variable Is Equal to One if Subject Chose with His Signal and Zero Otherwise

	Logistic regression		
Dependent variable: choose with signal	(1)	(2)	(3)
	All data	All data	First condition only
Position	−0.139***	−0.142***	−0.195***
	(0.051)	(0.050)	(0.062)
NPI	0.779***	0.776***	0.715***
	(0.121)	(0.120)	(0.114)
RT condition	−0.002	0.035	1.061
	(0.503)	(0.517)	(1.353)
Previous RT	0.051	0.054	0.174
	(0.109)	(0.110)	(0.245)
No signal match	−0.407	−0.420	−0.502
	(0.515)	(0.511)	(1.090)
No signal match × previous RT	0.076	−0.070	−0.086
	(0.122)	(0.119)	(0.257)
No signal match × RT condition	−0.677	−0.729	−1.664
	(0.512)	(0.534)	(1.358)
Previous RT × RT condition	−0.116	−0.121	−0.318
	(0.120)	(0.119)	(0.278)
No signal match × previous RT × RT condition	0.412***	0.355**	0.591**
	(0.146)	(0.155)	(0.296)
No-RT condition first		−0.176	
		(0.309)	
No-RT condition first × no signal match × previous RT × RT condition		0.171	
		(0.143)	
Constant	2.912***	2.980***	2.433**
	(0.555)	(0.523)	(1.161)
Observations	1,890	1,890	945
Pseudo-R^2	0.487	0.489	0.526

Notes. The variable *position* is the subject's position in the round (two to eight), *NPI* is net public information as defined in the text, *RT condition* is a dummy that takes on the value of one if the subject was in the RT condition, *previous RT* is the predecessor's RT (in seconds), *no signal match* takes on the value of one if the subject's private signal does not match his predecessor's move. *no-RT condition first* is a dummy that takes on the value of one if the subject was in a session where the no-RT condition was run first. Standard errors are clustered at the subject level and are shown in parentheses.

**Statistical significance at the 5% level; ***statistical significance at the 1% level.

found a small but insignificant increase in payoffs in the RT condition compared with the no-RT condition (\$0.58 versus \$0.56, $p = 0.56$).

Given the small predicted payoff effect in the previous analysis, we sought a more sensitive test of efficiency. In our simulations, most of the increase in efficiency comes from rounds in which incorrect cascades are triggered. In other words, when public information accurately reflects the true state of the world, there is little marginal value of RT information. To demonstrate this, we partition the simulated data into those rounds where an incorrect cascade occurs (the first two players receive incorrect signals) and compare the efficiency gains with all other rounds. Figure 7(b) summarizes the simulation results and displays two patterns. First, the average payoff is substantially lower for incorrect cascades compared with all other rounds (\$0.78 lower in the no-RT condition and \$0.32 lower in the RT condition). Second, and more importantly, the figure shows that the increase in efficiency because of RT information is positive (\$0.41) in an incorrect cascade, whereas it is slightly negative in all other rounds.

In Figure 7(c), we examine whether our experimental data are consistent with these two predictions. As expected, we do find that average payoffs are significantly lower in incorrect cascades (\$0.65 versus \$0.21, $p < 0.001$). Second, when conditioning on incorrect cascades, we find a significant increase in efficiency in the RT condition compared with the no-RT condition (\$0.30 versus \$0.11, $p = 0.008$).

Figure 6. (Color online) Probability of Choosing in Line with One's Private Signal as a Function of Condition and the RT of One's Predecessor

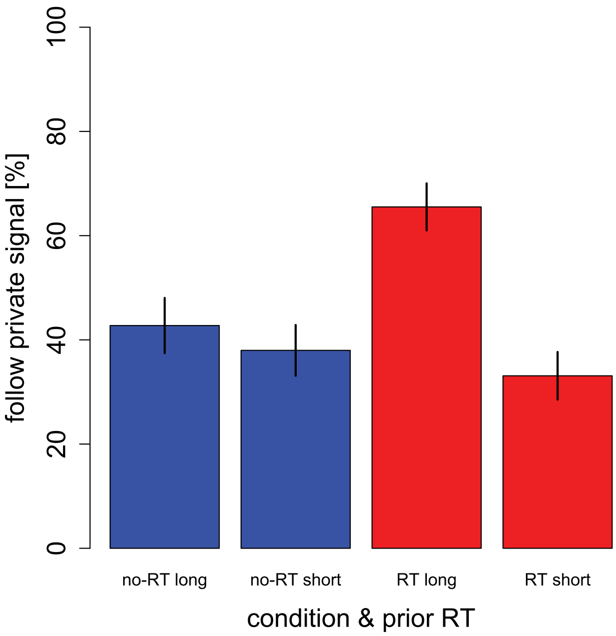

Notes. RTs are long/short if they are above/below the median RT of 2.03 seconds. In the RT condition, a subject was more likely to follow his private signal if his predecessor's RT was long rather than short. Data are restricted to positions 4–8 and conditioning on trials in which the private signal does not match the previous move. Vertical lines are standard errors of the mean, clustered by subject.

In summary, we find that RT information does generate an increase in efficiency, although this benefit is primarily concentrated in situations that occur infrequently by construction. Specifically, the probability of triggering an incorrect cascade equals the probability that the first two players both receive an incorrect signal: $(1-\frac{2}{3})^2$. This provides some intuition for the theoretical observation that cascades with mistaken beliefs arise only rarely.

5.4. Robustness Checks

In this section, we present a series of robustness checks on our main behavioral result that subjects are able to invert the mapping from private information to RTs and then, use this information in their decision.

5.4.1. Alternative Interpretation for Slow Decisions: Increased Effort. We have demonstrated that the RTs in our experiment are consistent with a key prediction of SSMs: that longer RTs are associated with indifference. However, as discussed in Section 3, longer RTs can also potentially reflect more effort on the part of the decision maker. This force could counteract the association between RT and indifference and makes it less valuable to discount a predecessor's slow decision.

If slow decisions are in fact driven more by effort than by conflict, then we would expect to observe that subjects attach more weight to their predecessors' slow decisions compared with fast decisions. Instead, the data show that subjects attach less weight to slow decisions (Figure 6).

Additionally, if slow decisions are driven more by effort than by conflict, then these decisions should, on average, be rewarded with higher payoffs. To investigate which force has the dominant effect on RTs, we ran a simple mixed-effects logistic regression of *correct choice* on *RT*. This regression reveals a strong negative effect in both the whole data set ($p = 10^{-6}$) as well as only in the no-RT condition ($p = 0.015$).

Next, we attempted to control for choice difficulty in the regression, by additionally including *NPI* and $|\,NPI\,|$. Both variables had significant positive effects on accuracy in the whole data set ($p = 10^{-12}$ and $p = 0.0005$, respectively) and in the no-RT condition ($p = 10^{-7}$ and $p = 0.004$, respectively). However, RT continued to have a negative, albeit insignificant, effect on accuracy in both the whole data set ($p = 0.11$) and no-RT condition ($p = 0.82$).

An alternative way to control for the effect of conflict is to average over all of a subject's decisions and test whether slower subjects are more accurate, on average. This analysis helps test the possibility that some subjects engage in more effortful thinking than other subjects. To test this, we correlated subject-level measures of accuracy with mean RT ($r = 0.23$, $p = 0.85$) and median RT ($r = -0.0035$, $p = 0.98$), the latter of which is shown in Figure A.3 in the online appendix.

Of course, we cannot be certain that we have completely accounted for conflict in these analyses, but they suggest that conflict plays a larger role than effort in generating RTs. In summary, although both forces are potentially at work in our experiment, the net result is that longer RTs do not signal higher accuracy.[11]

5.4.2. Experimenter Demand Effects. One potential concern with our design is that both our experimental instructions and the display screen make the RT information salient to subjects. Thus, it is possible that the magnitude of our RT effects on behavior is inflated. However, note that neither the experimental instructions nor the display tables provided any guidance on how to *interpret* the RT information (Figure 1). In other words, the behavioral effects we document are consistent with the class of SSMs described in Section 3, but subjects were not given any information about these models or their basic prediction that difficult decisions are, on average, slower.

Figure 7. (Color online) Efficiency Differences Across Experimental Conditions

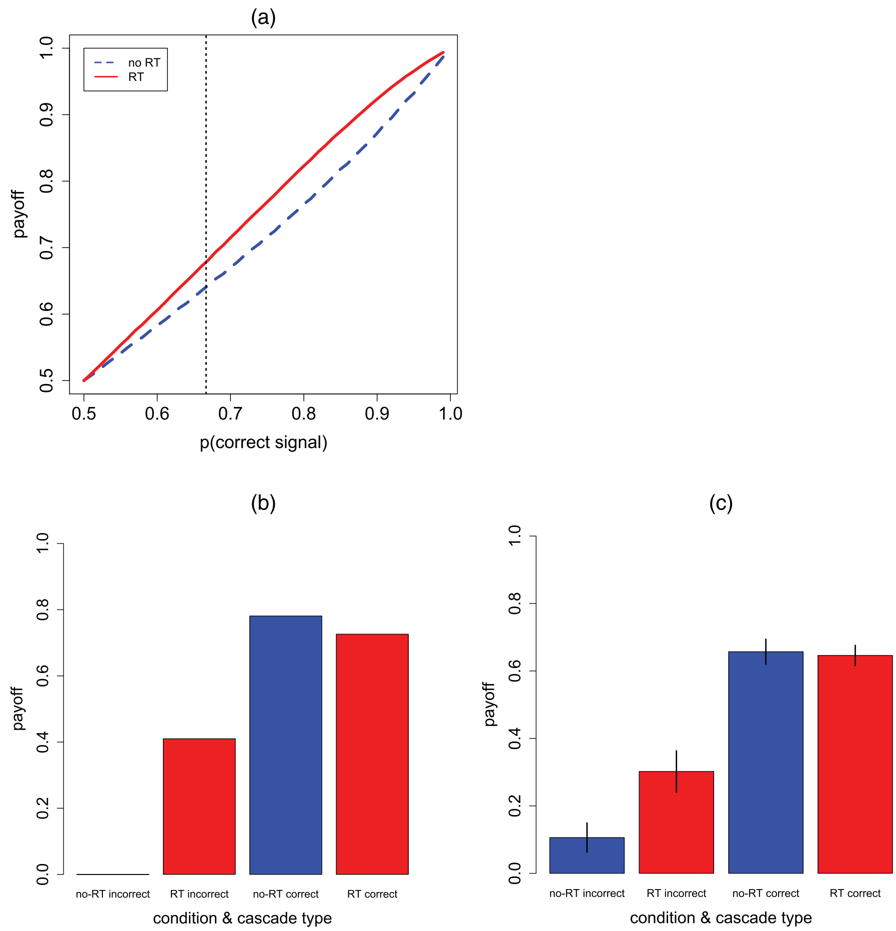

Notes. (a) Theoretical prediction of BNE with and without RT information. The average payoff is shown as a function of signal informativeness (our design used a value of 2/3, denoted by the vertical dotted line.). The average payoff is uniformly higher in the RT condition. (b and c) Average payoffs disaggregated by condition. "Incorrect" is defined as a round in which the first two subjects receive incorrect signals. "Correct" is defined as a round in which at least one of the first two subjects receives a correct signal. Panel (b) shows the theoretically predicted payoffs, and panel (c) shows the experimental data. Both simulations and experimental data are restricted to positions 4–8 and conditioning on trials in which the private signal does not match the previous move.

In our data, we can also partially address the experimenter demand concern by analyzing subsets of our data. In particular, the experimenter demand concern is most severe when analyzing RT effects within subjects. We can "transform" our design into a between-subjects design by exclusively analyzing data from the first half of each experimental session. The data we analyze are therefore generated from subjects who have only been exposed to one condition. Column (3) of Table 1 provides regression results using data only from the first 15 rounds of each experimental session. We find that the coefficient on the triple interaction remains significantly positive, which helps alleviate the within-subject design concern of the experimenter demand effect.

It is also worth noting that in our experiment, decisions are relatively fast compared with high-stakes decisions made in the field. For example, political endorsements or corporate mergers can take place over timescales of weeks or months, in which case RT information is easily measured and publicly available, as it is in our experiment (Giglio and Shue 2014). In addition, there are other potentially valuable features of the choice process that can be observed in the field—such as body language or tone of voice—which are shut down in our experiment by design. These process data can, in principle, be combined with RT information to further infer private information (Mayew and Venkatachalam 2012). As mentioned before, our design also mimics many online marketplaces where timing data are typically displayed next to each transaction so that agents who enter the market can easily observe its history.

6. Response Times in Economic Environments Outside of the Laboratory

One advantage of conducting our tests in the laboratory is that we can precisely control the provision of RT information to all players in the game. In this section, we discuss social learning settings outside the laboratory where timing information is also readily available. Importantly, the examples also retain the feature that actions are sequential and discrete, and thus, the availability of timing information can impact behavior. We begin with several examples from Bikhchandani et al. (1992).

6.1. Politics

Bikhchandani et al. (1992) discuss sequential polls and primaries and how they can generate "political momentum." Bartels (1988) argues that individuals' beliefs about a candidate can be heavily influenced by the decisions of others. More recent research has studied how the speed of endorsements for a political candidate can provide information about the strength of the endorser's preference for the candidate. Cohen et al. (2009) argue that slow endorsements indicate weak support and are less convincing compared with endorsements that come earlier in a political campaign.

Models and experiments on sequential voting have a similar flavor to the information cascades framework—except that a voter's payoff will depend on the actions of the entire electorate (Dekel and Piccione 2000, Battaglini et al. 2007, Callander 2007). When sequential votes are publicly observable, RT may have the capability to further aggregate private information (e.g., votes at city council meetings or roll call votes in U.S. national legislatures).

6.2. Finance

There are a variety of applications in financial markets where social learning is important: stock analyst forecasting, initial public offerings (IPOs), and information revelation in asset markets. With IPOs, a quote from Bikhchandani et al. (1992) is telling: "Welch (1992) uses a cascade model to show that if sufficiently many (few) individuals sign up *early* to receive shares, all (no) subsequent individuals follow their lead" (emphasis ours), p. 1013. It is thus clear that the timing of purchases both is available and is an important piece of information for subsequent investors. Among stock analysts, Welch (2000) shows that buy/sell recommendations have a significant impact on the recommendations of the subsequent two analysts. Although Welch (2000) does not analyze timing information explicitly, the date of each analyst recommendation is typically public information and can thus be used by subsequent analysts and investors to assess the confidence of the recommendation.

The speed of asset market transactions provides another good example where timing is useful in inferring information. Camerer and Weigelt (1991) and Bossaerts et al. (2014) provide experimental evidence that traders learn whether there are insiders in the market by observing the speed of trades (see also De Martino et al. 2013).

6.3. Peer Influence and Stigma

Bikhchandani et al. (1992) also discuss social phenomena potentially attributable to information cascades, such as typecasting. They argue that stigmas are often learned by "observing the behavior of others such as parents" (in which case, RT is certainly observable), p. 1013. They also argue that "a job applicant who receives *early* job offers may become a 'star'" (emphasis ours), p. 1014. Finally, in both personal and professional relationships, people observe (and often, extensively analyze) the time it takes for another person to respond to email or phone calls.

6.4. Additional Examples

In the domain of public health, an important topic is whether individuals take vaccines once they become available. Suppose that each individual receives a private signal about the net benefit of the vaccine. Those individuals who decide to become vaccinated can send a positive signal to others—for example, through a visible bracelet (Karing 2021); the strength of this signal can then be amplified if the decision to become vaccinated is made quickly.

Other examples pertain to transactions that are increasingly being carried out online. The timing of such transactions is always available to the firm providing the service or product, and so at the very least, they have access to timing information that could reveal consumers' private information (e.g., each individual's time spent considering a particular product).[12]

There are of course other settings where RT might contain private information and help others to resolve uncertainty. In markets for one of a kind goods such as art, sellers may not be able to estimate the values of their goods. Starting with a very high price and then observing potential buyers' RTs as they evaluate the offer may allow sellers to learn about the value of their good, with slower rejections signaling a higher-value item than faster rejections. In a sequential job search, a candidate might start by applying to the most ambitious job available and work his way down the list of options. However, because applying for jobs (and remaining unemployed) is costly, the job candidate may use the speed with which he is rejected to infer his value on the job market and jump farther down the list after a series of fast rejections. Similarly, an academic trying to judge the importance of his new journal article may start with the highest-ranked journals but move farther down the rankings if his paper is repeatedly rejected quickly.

7. Conclusion

In this paper, we conducted an experiment to test (1) whether RTs contain information about a subject's private information and (2) whether subjects infer private information from these RTs in a manner consistent with sequential sampling models. We find broad support in favor of both hypotheses. First, the distribution of RTs and choices in the no-RT condition is consistent with the prediction of slow indifference, whereby more conflicted decisions are associated with longer RTs. Specifically, we find that RTs peak when there is a net imbalance of two previous actions against a subject's private signal, and on these trials, subjects choose with their signal approximately 50% of the time. This result is important because it was previously unclear if the prediction of slow indifference, which has so far primarily been tested in simple subjective preference problems, would extend to more complex settings.

From an econometrician's point of view, this result is also interesting because the RT data provide extra information about the subject's private signal that is not contained in the choice data alone.[13] This result therefore complements recent research that documents some advantages of using RT data to better infer preferences (Chabris et al. 2009, Alós-Ferrer et al. 2018, Clithero 2018, Konovalov and Krajbich 2019). From a theoretical perspective, the result that RTs contain additional information beyond choice outcomes also suggests that RT data could be used by agents to achieve greater efficiency. Indeed, we found that in those situations where the standard model of social learning predicts that agents will "herd" on the incorrect action, the provision of RT information increases efficiency.

One interesting aspect of our RT results is that subjects in our experiment had a fixed decision time window of 10 seconds, meaning that they could not proceed through the experiment any faster by choosing more quickly. Nevertheless, we observed the standard relationship between strength of preference and RT. This challenges some of the literature assuming that RTs are the result of speed-accuracy trade-offs. In particular, subjects could not increase their reward rate by moving through the experiment more quickly; thus, the optimal strategy would have been to use the full 10 seconds to decide. On the other hand, if one thinks more broadly about rewards obtained outside of the experimental context, one can rationalize such behavior by assuming that subjects prefer to quickly execute a decision, in order to reduce mental effort (Kool and Botvinick 2014).

As a first pass at testing whether subjects use RTs in their decisions, we designed the experiment in what we believe to be the simplest possible setting. Neither subject i's RT nor the actions of subsequent agents affected his payoff. In this setting, there is no incentive for a subject to strategically transmit his private information. One could certainly imagine variants on this setting where a subject's payoff is a function of the accuracy of all agents in the group. In this case, subjects would have an incentive to strategically choose their action and RT. However, it is clear from the speed-accuracy trade-off that RTs cannot be manipulated arbitrarily without increasing the likelihood of an incorrect response. Thus, whether people attempt to manipulate their RTs is an interesting research direction. For instance, Konovalov and Krajbich (2017) investigate RT in two-stage bargaining games. They find that sellers can use buyers' initial rejection times to price discriminate against buyers with their second offers, earning higher profits. However, buyers who are made aware of this possibility strategically attempt to reject initial offers more quickly, in order to get lower second offers.

It is worth mentioning how our results relate to some other work on RT in economics. In particular, Eliaz and Rubinstein (2014) have modeled situations in which agents infer that decisions with long RTs are based on more deliberation and so, are more informative. In contrast, here we have argued that long RTs are typically indicative of indifference or conflicting information, and so, the resulting decision outcomes should be viewed with skepticism. To put it another way, RTs have often been treated as exogenous to the choice problem and primarily driven by the effort chosen by the agent. Instead, we argue that agents invest consistent levels of effort (i.e., setting evidence thresholds in the SSM), and the sources of variation in RTs are the endogenous features of the choice problems, which we capture by NPI (Moffatt 2005, Chabris et al. 2009, Dickhaut et al. 2009, Krajbich et al. 2014, Woodford 2014, Fudenberg et al. 2018, Konovalov and Krajbich 2019, Schotter and Trevino 2020). Consistent with this interpretation, we find no evidence that slower subjects are more accurate in our experiment. In future work, it will be important to study the distinguishing features of these two types of decision environments.

We also note that there is some precedent for using RT as a measure of beliefs in experimental economics. For example, Frydman and Nave (2017) demonstrate that RTs correlate with beliefs in a perceptual discrimination task, and moreover, these RTs contain information about belief updating in a separate economic context. Frydman and Nave (2017) also employ an SSM to estimate trial by trial belief updating from RT data. Turocy and Cason (2015) demonstrate that RT can be used to infer private information in experimental auctions. Their data indicate that RT increases with a bidder's private signal in first-price auctions, but not for second-price auctions. There is also evidence that RT can be used to classify subjects into sophisticated versus naïve "types" (Agranov et al. 2015). Taken together, these papers suggest that RTs are informative about beliefs in economic experiments. This aligns well with our results because beliefs are presumably a function of public and private information in cascade games. We add to this literature by additionally testing whether subjects—and not just the researcher—are able to extract information contained in RT.

Finally, the decisions in our experiment were made relatively quickly, compared with what we would expect in a high-stakes, single-shot decision. There is a potential worry that subjects might therefore be relying on simple heuristics that would not apply in "real-world" situations. However, the use of simple heuristics would only hinder our (and the subjects') ability to infer information from the RTs because it would make responses more automatic and less sensitive to the weight of information. In high-stakes decisions, we would only expect RTs to be *more* informative. In such settings, we would expect agents to increase emphasis on accuracy at the cost of speed. This in turn would generate a larger gap between average RTs for easy versus difficult problems. It is also important to recognize that many economic decisions are made quickly (e.g., choosing between a high-quality, high-price versus low-quality, low-price item at the grocery store), which suggests that slow indifference may occur in a substantial portion of economic transactions. Although our paper provides evidence that RT plays an important role in a stylized information cascade environment, more research is needed to explore RT in more general economic environments.

Acknowledgments

The authors thank Chris Crabbe for programming the experiment and Xiao Luo, Naomi Anderson, Cameron Luther, and Constantin Charles for research assistance. They also thank James Peck and Salvatore Nunnari for comments on the manuscript. For helpful discussions, they thank Colin Camerer, Chad Kendall, Todd Hare, Michael Alvarez, Paul J. Healy, Muriel Niederle, and Katie Coffman.

Endnotes

[1] Van de Calseyde et al. (2014) also provide further evidence from the field using data from the television show *The Voice*, where contestants perform an audition in front of a set of competing coaches. After the audition begins, coaches have two minutes to indicate whether they want to coach the candidate. The authors find that when two coaches both indicate they would like to coach the candidate, 70% of candidates choose the coach with the shortest response time.

[2] Although the information cascade literature is often invoked to explain inefficiencies associated with social learning, there are very restrictive conditions under which an agent will form a confident yet incorrect belief. As Eyster and Rabin (2010) put it, "the rarity with which rational learning produces strong mistaken beliefs is perhaps insufficiently salient in economists' conception of the rational-herding literature. Rational herders either converge to only weak public beliefs or only very infrequently herd on the wrong action" (Eyster and Rabin 2010, p. 226).

[3] See also the work by Cipriani and Guardino (2005) and Drehmann et al. (2005), who conduct experimental tests of social learning in a financial market setting (Avery and Zemsky 1998).

[4] For examples from an individual decision-making environment, see Wilcox (1993); Moffatt (2005); Dickhaut et al. (2009); Achtziger and Alós-Ferrer (2013); Paravisini and Schoar (2013); Krajbich et al. (2014); Fudenberg et al. (2018); Alós-Ferrer et al. (2016); Rubinstein (2016); Agranov and Ortoleva (2017); Echenique and Saito (2017); Spiliopoulos and Ortmann (2017); Clithero (2018); and Webb (2018). For examples from a strategic environment, see Agranov et al. (2015), Gill and Prowse (2017), Konovalov and Krajbich (2020), and Schotter and Trevino (2020).

[5] For an alternative experimental design, where subjects were not required (or allowed) to press the "continue" button after receiving their private signal, see the online appendix.

[6] We chose the 10-second cutoff to give subjects enough time to make their decisions without feeling rushed. Moreover, this design choice was based on pilot testing where we observed that implementing a more severe time limit would likely begin to constrain some subjects.

[7] The unit of observation in the logistic regression is at the subject-decision level. Unless otherwise specified, all subsequent analyses reported in the paper are also at the subject-decision level.

[8] Subjects' behavior over the course of the experiment is fairly stable. When pooling all decisions across both experimental conditions, we found that the probability of choosing with their signal increases modestly over the 30 rounds ($p = 0.07$) (see Figure A.1 in the online appendix), although there is no significant difference when analyzing this trend within each of the two experimental conditions.

[9] We use the median rather than the mean because typically, RTs are approximately log-normally distributed (as they are in the current data set). Indeed, this is a key feature of RTs generated by SSMs, and thus, the median provides a more appropriate summary statistic given the positively skewed distribution.

[10] The median is computed using the distribution of predecessor RTs from trials where subjects received a different signal than their predecessor and were located after the third position. Our conclusion does not change if we use an alternative definition of the median based on the distribution of RTs in the entire data set (1,888 ms).

[11] A distinct alternative explanation is that trials with slow RTs are those in which subjects are distracted by other activities. However, this hypothesis cannot explain the full extent of the behavior: Figure 3 demonstrates that RTs systematically rise as NPI approaches −2, and there is no reason that distraction should correlate with NPI.

[12] There are several other applied settings where response times have been shown to contain information about economic variables, although not necessarily in a strategic environment. For example, Paravisini and Schoar (2013) show that response times contain information about credit applications. They find that the length of time (typically between 2 and 10 minutes) that a loan officer takes to evaluate a loan application is increasing in the requested loan amount. In a corporate finance setting, Giglio and Shue (2014) show that the passage of time contains information about the probability a merger will be completed.

[13] Our methodology also provides an alternative to other belief elicitation procedures in information cascade experiments (Penczynksi 2017).

References

Achtziger A, Alós-Ferrer C (2013) Fast or rational? A response-times study of Bayesian updating. *Management Sci.* 60(4):923–938.

Agranov M, Ortoleva P (2017) Stochastic choice and preferences for randomization. *J. Political Econom.* 12(1):40–68.

Agranov M, Caplin A, Tergiman C (2015) Naive play and the process of choice in guessing games. *J. Econom. Sci. Assoc.* 1(2):146–157.

Alós-Ferrer C, Fehr E, Netzer N (2018) Time will tell: Recovering preferences when choices are noisy. CESifo Working Paper Series No. 7333, CESifo Group, Munich, Germany.

Alós-Ferrer C, Granić -G, Kern J, Wagner AK (2016) Preference reversals: Time and again. *J. Risk Uncertainity* 52(1):65–97.

Anderson LR, Holt CA (1997) Information cascades in the laboratory. *Amer. Econom. Rev.* 87(5):847–862.

Avery C, Zemsky P (1998) Multidimensional uncertainty and herd behavior in financial markets. *Amer. Econom. Rev.* 724–748.

Banerjee AV (1992) A simple model of herd behavior. *Quart. J. Econom.* 107(3):797–817.

Bartels L (1988) *Presidential Primaries and the Dynamics of Public Choice* (Princeton University Press, Princeton, NJ).

Battaglini M, Morton R, Palfrey T (2007) Efficiency, equity, and timing of voting mechanisms. *Amer. Political Sci. Rev.* 101(3):409–424.

Becker WJ, Connolly T, Slaughter JE (2010) The effect of job offer timing on offer acceptance, performance, and turnover. *Personnel Psych.* 63(1):223–241.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Political Econom*. 100(5):992–1026.

Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psych. Rev.* 113(4):700–765.

Bossaerts P, Frydman C, Ledyard J (2014) The speed of information revelation and eventual price quality in markets with insiders: Comparing two theories. *Rev. Finance* 18(1):1–22.

Busemeyer JR, Townsend JT (1993) Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psych. Rev.* 100(3):432–459.

Callander S (2007) Bandwagons and momentum in sequential voting. *Rev. Econom. Stud.* 74(3):653–684.

Camerer C, Weigelt K (1991) Information mirages in experimental asset markets. *J. Bus.* 64(4):463–493.

Camerer CF, Ho T-H, Chong J-K (2004) A cognitive hierarchy model of games. *Quart. J. Econom.* 119(3):861–898.

Chabris CF, Morris CL, Taubinsky D, Laibson D, Schuldt JP (2009) The allocation of time in decision-making. *J. Eur. Econom. Assoc.* 7(2–3):628–637.

Chamley C, Gale D (1994) Information revelation and strategic delay in a model of investment. *Econometrica* 1065–1085.

Cipriani M, Guarino A (2005) Herd behavior in a laboratory financial market. *Amer. Econom. Rev*. 95(5):1427–1443.

Clithero JA (2018) Improving out-of-sample predictions using response times and a model of the decision process. *J. Econom. Behav. Organ.* 148(April):344–375.

Cohen M, Karol D, Noel H, Zaller J (2009) *The Party Decides: Presidential Nominations Before and After Reform* (University of Chicago Press, Chicago).

Coricelli G, Nagel R (2009) Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proc. Natl. Acad. Sci. USA* 106(23):9163–9168.

Crawford VP, Iriberri N (2007) Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75(6):1721–1770.

Dekel E, Piccione M (2000) Sequential voting procedures in symmetric binary elections. *J. Political Econom.* 108(1):34–55.

De Martino B, O'Doherty JP, Ray D, Bossaerts P, Camerer C (2013) In the mind of the market: Theory of mind biases value computation during financial bubbles. *Neuron* 79(6):1222–1231.

Dickhaut J, Rustichini A, Smith V (2009) A neuroeconomic theory of the decision process. *Proc. Natl. Acad. Sci. USA* 106(52):22145–22150.

Drehmann M, Oechssler J, Roider A (2005) Herding and contrarian behavior in financial markets: An internet experiment. *Amer. Econom. Rev*. 95(5):1403–1426.

Echenique F, Saito K (2017) Response time and utility. *J. Econom. Behav. Organ.* 139(July):49–59.

Eliaz K, Rubinstein A (2014) A model of boundedly rational 'neuro' agents. *Econom. Theory* 57(3):515–528.

Eyster E, Rabin M (2010) Naïve herding in rich-information settings. *Amer. Econom. J. Microeconomics* 2(4):221–243.

Eyster E, Rabin M (2014) Extensive imitation is irrational and harmful. *Quart. J. Econom.* 129(4):1861–1898.

Eyster E, Rabin M, Weizsacker G (2015) An experiment on social mislearning. Preprint, submitted December 18, http://dx.doi.org/10.2139/ssrn.2704746.

Frydman C, Nave G (2017) Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism. *Management Sci.* 63(7):2340–2352.

Fudenberg D, Strack P, Strzalecki T (2018) Speed, accuracy, and the optimal timing of choices. *Amer. Econom. Rev.* 108(12):3651–3684.

Giglio S, Shue K (2014) No news is news: Do markets underreact to nothing? *Rev. Financial Stud.* 27(12):3389–3440.

Gill D, Prowse VL (2017) Using response times to measure strategic complexity and the value of thinking in games. Preprint, submitted January 19, http://dx.doi.org/10.2139/ssrn.2902411.

Goeree JK, Palfrey TR, Rogers BW, McKelvey RD (2007) Self-correcting information cascades. *Rev. Econom. Stud.* 74(3):733–762.

Guarino A, Harmgart H, Huck S (2011) Aggregate information cascades. *Games Econom. Behav*. 73(1):167–185.

Gul F, Lundholm R (1995) Endogenous timing and the clustering of agents' decisions. *J. Political Econom*. 103(5):1039–1066.

Herrera H, Hörner J (2013) Biased social learning. *Games Econom. Behav.* 80:131–146.

Karing A (2021) Social signaling and childhood immunization: A field experiment in Sierra Leone. Working paper, Princeton University.

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

Konovalov A, Krajbich I (2020) Decision times reveal private information in strategic settings: Evidence from bargaining experiments. Preprint, submitted September 22, http://dx.doi.org/10.2139/ssrn.3023640.

Kool W, Botvinick M (2014) A labor/leisure tradeoff in cognitive control. *J. Experiment. Psych. General* 143(1):131–141.

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.† *Amer. Econom. Rev.* 104(5):501–506.

Kübler D, Weizsäcker G (2004) Limited depth of reasoning and failure of cascade formation in the laboratory. *Rev. Econom. Stud.* 71(2):425–441.

Kübler D, Weizsäcker G (2005) Are longer cascades more stable? *J. Eur. Econom. Assoc.* 3(2–3):330–339.

Lee IH (1993) On the convergence of informational cascades. *J. Econom. Theory* 61(2):395–411.

Mayew WJ, Venkatachalam M (2012) The power of voice: Managerial affective states and future firm performance. *J. Finance* 67(1): 1–43.

McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal form games. *Games Econom. Behav.* 10(1):6–38.

Moffatt PG (2005) Stochastic choice and the allocation of cognitive effort. *Experiment. Econom.* 8(4):369–388.

Mormann MM, Malmaud J, Huth A, Koch C, Rangel A (2010) The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment Decision Making* 5(6):437–449.

Mosteller F, Nogee P (1951) An experimental measurement of utility. *J. Political Econom.* 59(5):371–404.

Paravisini D, Schoar A (2013) The incentive effect of scores: Randomized evidence from credit committees. NBER Working Paper No. 19303, National Bureau of Economic Research, Cambridge, MA.

Penczynski S (2017) The nature of social learning: Experimental evidence. *Eur. Econom. Rev.* 94:148–165.

Philiastides MG, Ratcliff R (2013) Influence of branding on preference-based decision making. *Psych. Sci.* 24(7):1208–1215.

Polanía R, Krajbich I, Grueschow M, Ruff CC (2014) Neural oscillations and synchronization differentially support evidence accumulation in perceptual and value-based decision making. *Neuron* 82(3):709–720.

Rodriguez CA, Turner BM, McClure SM (2014) Intertemporal choice as discounted value accumulation. *PLoS One* 9(2):e90138.

Rubinstein A (2016) A typology of players: Between instinctive and contemplative. *Quart. J. Econom.* 131(2):859–890.

Schotter A, Trevino I (2020) Is response time predictive of choice? An experimental study of threshold strategies. *Experiment. Econom.*, ePub ahead of print March 12, https://doi.org/10.1007/s10683-020-09651-1.

Spiliopoulos L, Ortmann A (2017) The BCD of response time analysis in experimental economics. *Experiment. Econom.* 21(2018): 383–433.

Stahl II DO, Wilson PW (1994) Experimental evidence on players' models of other players. *J. Econom. Behav. Organ.* 25(3):309–327.

Turocy TL, Cason TN (2015) Bidding in first-price and second-price interdependent-values auctions: A laboratory experiment. Centre for Behavioural and Experimental Social Science Working Paper No. 15-23, School of Economics, University of East Anglia, Norwich, UK.

Van de Calseyde PP, Keren G, Zeelenberg M (2014) Decision time as information in judgment and choice. *Organ. Behav. Human Decision Processes* 125(2):113–122.

Wald A (1945) Sequential tests of statistical hypotheses. *Ann. Math. Statist.* 16(2):117–186.

Webb R (2018) The (neural) dynamics of stochastic choice. *Management Sci.* 65(1):230–255.

Weizsäcker G (2010) Do we follow others when we should? A simple test of rational expectations. *Amer. Econom. Rev.* 100(5): 2340–2360.

Welch I (1992) Sequential sales, learning, and cascades. *J. Finance* 47(2):695–732.

Welch I (2000) Herding among security analysts. *J. Financial Econom.* 58(3):369–396.

Wilcox NT (1993) Lottery choice: Incentives, complexity and decision time. *Econom. J.* 103(421):1397–1417.

Woodford M (2014) Stochastic choice: An optimizing neuroeconomic model. *Amer. Econom. Rev.* 104(5):495–500.

Zhang J (1997) Strategic delay and the onset of investment cascades. *RAND J. Econom.* 28(1):188–205.