Risk-Seeking in a Continuous Game of Timing

David Reitter, Jens Grossklags, Alan Nochenson
reitter@psu.edu, jensg@ist.psu.edu, aen137@ist.psu.edu
College of Information Sciences and Technology
The Pennsylvania State University

Abstract

Continuous games of timing represent an ecologically relevant case of decision-making under uncertainty with many applications to tactical security scenarios. We present two experiments collecting data using the two-player, 20-second FlipIt game, in which a player has to react to an estimated, externally controlled quantity in light of a continuous payoff distribution. The quantity is represented once as time interval, and once as distance in time and space. Risk-taking propensity is an individual parameter obtained via a standardized survey instrument. We find that risk propensity affects participants' decisions in early and late rounds of the game and show that risk-seeking individuals fail to play well initially in the difficult, time-interval estimation task, but that they can be trained to make rational choices. Participants' biases are gaussian-distributed. They do not reflect the distribution of payoffs or expected utility. We present an ACT-R model using established error distributions for temporal and visual perception and instance-based learning.

Keywords: Decision-making under uncertainty; Risk propensity; Games of timing; Security economics; ACT-R

Introduction

Decision-making in the real world is frequently dynamic. Choices are made in an incremental fashion, and in continuous time. Further, our decisions are often not aligned with other individuals' actions (i.e., the choices are made in a sequential or asynchronous fashion). Unsurprisingly, acquiring and processing feedback as well as developing a sound strategy in such environments is particularly challenging (Friedman, Shor, Shenker, & Sopher, 2004; Weber & Johnson, 2009).

Further, being successful in dynamic environments shifts the focus from selecting the most suitable option from a pool of alternatives to a decision problem of *when* to act to get an advantage over an opponent. For example, in tactical security scenarios it is important to jump to action at the right time to avoid a loss of money or even human life. In such *games of timing*, individuals have to manage the risk of acting too early or too late (i.e., they have to minimize their individual bias with respect to timing).

In our work, we behaviorally study and model a specific game of timing, the FlipIt game, in which one human and one computerized player compete for ownership of a game board (van Dijk, Juels, Oprea, & Rivest, 2012). We pursue two objectives with our analysis. First, we want to understand how the modality of presentation of real-time stimuli influences performance in the game. We distinguish between subexperiments that either allow the estimation of the opponent's timing from permanently available information, or force the individual to infer the same information from transient visualizations. Second, we expand our analysis to investigate how individual differences in risk perception are related to an individual's bias in making timely choices. For this purpose, we utilize a survey instrument, the Risk Propensity Scale, which measures an individual's general tendency towards risk-taking in everyday tasks (Meertens & Lion, 2008).

We thoroughly analyze the data of 230 individuals who participated in a FlipIt game with performance-based compensation. With the insights from the experimental data analysis, we develop an ACT-R model using instance-based learning and established error distributions for the perception of the different stimuli. We expect that such models will ultimately integrate quantifiable, computational accounts of perception and memory with those of individual differences in risk-taking and economic decision-making under uncertainty.

Related Work

Games of Timing in Theory and Experiment: Non-cooperative games with continuous timing and asynchronous decision-making have been investigated with formal methods since the cold war era (e.g., Blackwell, 1949; Radzik & Orlowski, 1982). These theoretical works have been complemented with behavioral and experimental studies (see, e.g., Kahan & Rapoport, 1974; Rapoport & Murphy, 2012). More recently, Brunnermeier and Morgan studied multi-player games where agents receive private signals about a payoff-relevant state variable in a scenario with individualized desynchronized clocks. To perform well in the game, agents have to predict other agents' clock times subject to different information conditions and the number of players (Brunnermeier & Morgan, 2010).

Continuous time decision-making is challenging for humans and game-theoretical predictions frequently fail to explain experimental observations. For example, Friedman et al. observe that convergence can fail even when iterated deletion of dominated strategies would theoretically lead to the Nash equilibrium (Friedman et al., 2004).

Decision-Making under Risk: Research on decision-making under risk has evolved from normative axiomatic models toward empirically-supported frameworks. This shift has been motivated by field and laboratory experiments showing that actual behaviors contradict even the most basic axiomatic assumptions (see Weber & Johnson, 2009, for a comprehensive survey). We want to highlight two primary themes that have emerged.

First, general theories have been proposed that predict behavior across different canonical situations. Most prominently, Prospect Theory suggests that individuals will be risk averse in the domain of gains, and risk seeking in the domain of losses (Kahneman & Tversky, 1979). The theory, for example, explains experimental observations in which a simple reframing of a choice situation yields different human decisions even if the available information is equivalent. Second, experimental studies have isolated a number of individual difference variables (e.g., personality traits) which have shown to impact decision-making under risk. For example, Zuckerman et al. showed that the propensity for sensation seeking is related to the involvement in high-risk activities (Zuckerman, Kolin, Price, & Zoob, 1964).

Recent work has focused on developing measurement constructs to evaluate individuals' general risk taking. Suitable to our context, we selected a scale which measures everyday risk-taking behavior, the Risk Propensity Scale (Meertens & Lion, 2008). Applying this scale to scenarios with different information availability, it has been shown that risk propensity primarily influences decisions in which individuals face an unknown risk (Meertens & Lion, 2011).

We are also interested in the impact of learning on risk attitudes. Concerning this question, March proposes that when individuals are presented with options with positive outcomes, then learning will lead them to favor safer options in the future (March, 1996). He argues that the opposite holds for learning in the negative domain (i.e., they would converge to more risky options). However, he cautions that the latter effect might be short-lived and individuals tend to migrate to more risk neutral choices in the long run. In a follow-up economic model, this observation is explained with the help of an motivation-driven convergence process (DellaVigna & LiCalzi, 2001). As a result, learning adjusts an individual's reference point leading to risk neutral choices while her risk attitudes remain consistent with Prospect Theory.

Experimental Method

The Game: FlipIt is an experimental game for two players, *red* and *blue*, who vie for ownership of a common gameboard by calling for a "flip" (van Dijk et al., 2012). The goal of each player is to "own" the gameboard for as much time as possible. Each experiment consists of 6 rounds that are relatively fast-paced with each round lasting 20 seconds. Players gain 100 points for every second they own the board. Flips are costly, i.e., each time a player calls for a flip, they are charged 100 points. In this version of the game, *blue* is a human player, while *red* is the computer opponent. The game is symmetric.

Without further restrictions, the rational behavior for a player would be to flip the board back as soon as possible after the opponent has flipped. Anticipating the opponent to be rational, maintaining the status quo and not flipping at all would be the rational response (a recursive game, e.g., Traveller's Dillemma).

However, FlipIt does not make information fully available to the player. The player will not see an opponent's flip in real time and therefore must estimate when to expect it. Once the player has flipped the board, he can see how long ago the opponent flipped, if so. If a player calls for a flip too early, they are still charged for the flip, but do not gain anything, as they already own the board. Thus, their payoff is optimal immediately after an opponent's flip, decreases linearly afterwards, but is negative shortly before the opponent's flip.

The computer plays in a predictable manner: in our experiment, it attempts to flip at constant intervals. So, assume a subject can arrive at a noisy estimate of the optimal time to act; the estimate will improve over time as more data becomes available. The FlipIt game used here distributes the payoff such that early action is penalized severely, but late action is acceptable (see subfigure in Fig. 3). To manage risk, the human player has to evaluate the reliability of his estimate and take into account the payoff distribution.

In game-theoretic experiments, payoff matrices are typically stated explicitly to the subject; decisions are thus thought to be a combination of explicit, System-2 reasoning, and implicit, subsymbolic cognition. Are payoff matrices a good model of rapid System-1 decision-making? What if the payoff distribution is smooth, and discontinuous, as in our timing game?

This question can be explored, if not definitively tested, in a short experiment, in which the payoff function is not directly given to subjects, even though the game is thoroughly explained.

If subsymbolic sensitivity to underlying payoff functions exists, then we expect people to be sensitive to the payoff distribution around the estimate, and to modify their bias accordingly. Their risk propensity personality trait would interact with their bias. (Bias, here, is used to describe any deviation from the maximum-likelihood estimate to account for risk.)

Strategy of the Computer Player: Each participant faced an opponent playing a fixed (non-adaptive) periodic strategy. The opponent's flip rate and time of first flip (anchor, or phase) changed in every round of the game, but the overall strategy of the opponent did not. Both values were drawn from uniform distributions for each round. Flip rates ranged from 1 to 5 seconds, and the anchor ranged from 0.1 to 4.1 seconds.

Experiments: We report two experiments: the first, with permanent visual information availability, the second with only transient availability of information. A total of 400 subjects based in the U.S.A. participated for compensation on the Amazon Mechanical Turk platform. Data from 46 participants were excluded for failing to complete all rounds or for double participation due to a computer error.[1]

Experiment 1: An example of the feedback shown to players (at the end of a round) in this experiment is shown in Figure 1. 203 subjects completed this experiment.

Experiment 2: The first experiment allowed subjects to use the full flip history of the round to determine the opponent's

[1]246 of the included participants were males. The mean age of all participants was 29.3 $sd = 10.2$. A superset of the Exp. 1 data is reported in (Nochenson & Grossklags, 2013); here we report results on a subset that includes only the information conditions used in both experiments.

Figure 1: Feedback shown in Experiment 1. This information was updated after each flip, so that participants saw data to the left of their last flip. Also shown was a progressing grey bar to the right. Blue dots represent flips by the participant, red dots indicate opponent flips.

Figure 2: Feedback shown in Experiment 2. This feedback was shown after each "successful" flip that gave the player back control. Blue area is area that was under control by the participant until the red player took control. Each red X indicates that the opponent made a superfluous flip (the position of each X is inconsequential).

rate of flip; the history was presented visually. This modality lends itself to graphical reasoning; it may not be representative of the security situations that motivated the work. In the second experiment, the intention was to keep the task, the total available information and hence the explicit reasoning demanded of a subject the same. However, we now present information in a transient manner, requiring the player to produce estimates of the opponent behavior from memory.

151 participants completed the second experiment. An example of the feedback shown to participants in the second experiment (after each of their flips) is shown in Figure 2.

In each experiment, we systematically varied the amount of *a priori* information that was given to participants about the strategy of their opponent. We varied this information in order to broaden the range of analytical thinking that may be taking place depending on the amount of information given (see Nochenson & Grossklags, 2013, for an analysis). Treatment assignment was constant across all rounds for a given participant, and controlled across both experiments.

Procedure, Survey and Payments: Participants received detailed written instructions with graphical representations of the game (adjusted for each type of experiment), a survey comprising 4 demographic questions, 7 questions assessing risk propensity (Meertens & Lion, 2008), 5 questions assessing need for cognition (Wood & Swait, 2002), and 3 basic integrity questions ("screeners").

Following this, participants completed one unpaid round, and five paid rounds. Participants received a show-up fee (USD \$0.50) and a performance-based compensation (\$0.10 – \$1.27). At equal performance with the opponent, subjects were paid \$0.10 per round (i.e., \$0.50 for break-even performance across the five compensated rounds).

Experimental Results

In the following analysis, data points represent individual subjects' flips. The subject population and the fact that we investigate early-stage, untrained behavior necessitates more filtering for motivated and able subjects than would be adequate in a controlled laboratory setting. Therefore, we will focus on the subset of subjects ($n = 230$) that performed better than the simplistic computer opponent.[2]

We observe two dependent measures. The *round payment* reflects a subject's bonus payment per round. We see it as indicative of comprehending the overall game and of performing well at the core risk management and estimation task. The *timing bias* quantifies performance at the more precise task of estimating the opponent's flip time and triggering a flip in relation to it. The timing bias is defined as the temporal distance between a subject's flip and the nearest opponent flip (negative values indicate that the subject flipped shortly before the opponent). The theoretically optimal subject flip occurs immediately after an opponent's flip (timing bias $0 + \varepsilon$). We normalize the error by the periodicity of the opponent's flips, so that $-0.5/+.5$ would describe a flip that occurred halfway between two opponent's flips.

Histograms of the timing bias for those subjects who earned at least a break-even bonus payment show a peak following 0. (Subjects who lost all or part of their bonus endowment of \$.50 exhibit a peak before 0, i.e., before the opponent flip.)

Risk propensity affects people's behavior in the FlipIt game (Figures 3 and 4). Low-risk propensity subjects show different behavior in Exp. 1, where the full flip history is available visually: they do not concentrate their flips in the optimal zone. High-risk propensity participants, on the contrary, exhibit a peak shortly after the opponent's flip. The cognitive model will explain this as the result of subjects more reliably predicting the opponent's future flips. In Exp. 2, both participant groups show the peak at the correct time, although high-risk participants concentrate their flips near the opponent's flip, as would be expected.

Do people adapt to the underlying payoff structure of the experiment? For the qualitative payoff structure, the answer is "yes": participants understand the basic implications of the rules. However, for the quantitative payoff function (see subgraph in Figure 3), the answer is more complex. It appears that only individuals with high risk propensity show flip preferences congruent with payoff potential: they mostly flip shortly after the opponent, tapering off slowly afterwards, but strictly avoid the penalty zone right before the opponent's flip. Low-risk individuals show a symmetrical gaussian peak after the opponent's tick. The following explanation is consistent with the data: high-risk individuals are not just risk-takers,

[2]Similarly, our ACT-R model will focus on System-1 (implicit) learning and adaptation.

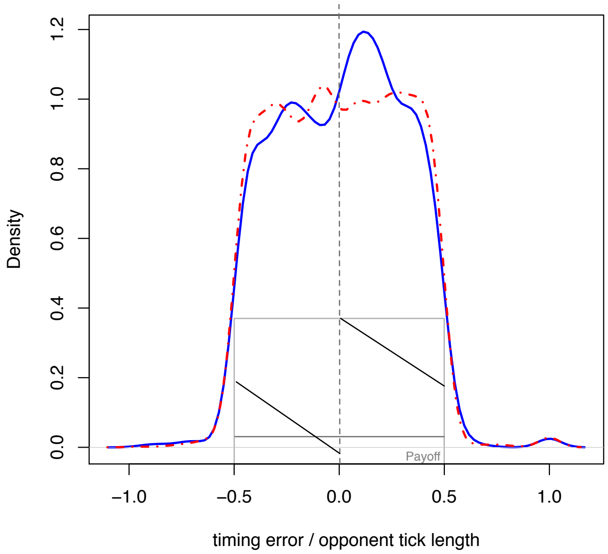

Figure 3: Distribution of timing bias in (empirical) experiments 1 (dotted) and 2 (solid), for **lower**-than-mean risk propensity personalities. The (scale-free) subgraph shows the discontinuous distribution of payoffs over time.

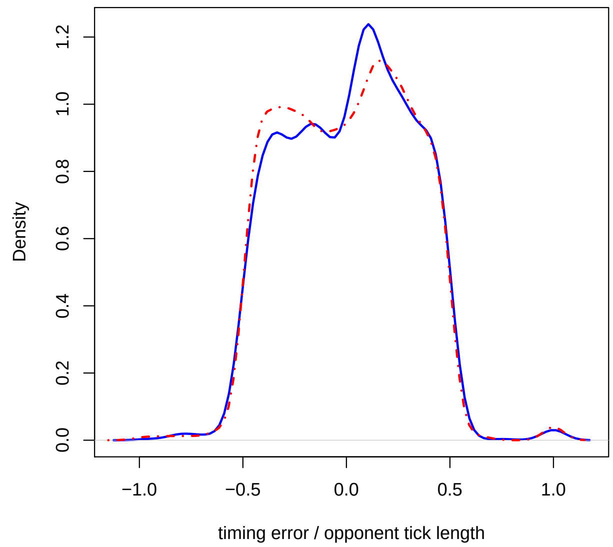

Figure 4: Timing bias as in Fig. 3, but for **higher**-than-mean risk propensity personalities.

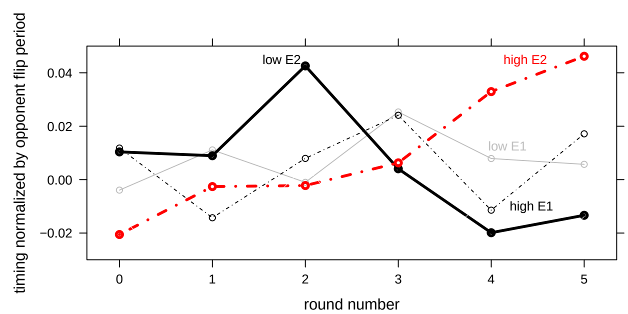

Figure 5: Timing for high vs low risp propensity individuals in both experiments at different levels of experience.

but also know better how to manage risk.

The distributions show an aggregate over all rounds, i.e., over all stages of learning. A linear mixed-effects regression model reveals a more detailed picture (see: suppl. mat., Table 1). Predicting the bonus payment gained by participants in each round shows that overall, in Exp. 1, low- and high-risk individuals fare similarly on average, but that high-risk propensity individuals achieve higher gains as time goes on.

Figure 5 visualizes the reasons for this effect: in Exp. 2, risk-seeking individuals start with an overly early timing (before the opponent's flip), but then shift their timing to be later as they learn to do the task. Risk-avoiding individuals begin with a low bias away from the flip and then adjust their timing as they learn. A linear-mixed effects model shows a significant, positive interaction effect (see: suppl. mat., Table 2) of experience and risk propensity in Exp. 2 on timing.

This provides empirical confirmation of DellaVigna and LiCalzi's hypothesis (see related work section) - although a qualified one: The modality of estimation (time vs. space) influences decision-making by risk-seekers. We propose that the mechanism of this interaction results from differences in how confident subjects are about their own estimates. If they are less confident about their abilities due to noisy estimates or the availability of information (as in Exp. 2), they will resort to their preferred default and seek higher or lower risks. Even short training can invert the effect due to a change in one's metacognitive assessment of task ability.

ACT-R Model

The cognitive model we propose first describes a developmental step that determines an individual's risk propensity R as the result of experience. Experience also allows the model to acquire a sense of error associated with different types of estimates.

The second component of the model implements decision-making strategies, separately for experiments 1 and 2. The model is situated within the ACT-R cognitive architecture, and is implemented using the ACT-UP library (Reitter & Lebiere, 2010). It makes use of Instance-Based Learning (Gonzalez, Lerch, & Lebiere, 2003), which allows the model to estimate an average value biased by recency and frequency from a series of observations.

Model 1 has the history of the opponent's flip behavior available visually. Thus, it does not need to memorize past flips, and instead uses eye movements over the horizontally spaced flip indicators on the screen. During each saccade, the distance is estimated. The model assumes a Gaussian distribution around the actual point as landing site (standard deviation of 0.1 distance in accordance with the ACT-R visual module, as in Salvucci, 2001). The resulting average estimate will be most precise when multiple opponent flips are present. A final saccade is made, spaced according to this estimate (with noise according to the above distribution). Thus, the model adjusts for risk using the RiskAdjustment procedure (described below), it foveates the resulting point and calls for a flip as soon as the time-indicator has reached it.

Model 2 relies on declarative memory to track the opponent's actions. It uses the ACT-R temporal module (Taatgen, van Rijn, & Anderson, 2007; Taatgen & van Rijn, 2011) to track temporal differences in events. Described with simplifications, the temporal module defined a periodic process (one tick per second), in which each tick is subject to noise (0.015 by default). Just like the visual model, this model assumes that the opponent's flips are periodic and very regular. It attempts to adopt a pace that is similar to the opponent's. To do so, it pays attention to feedback obtained after each flip: if the model's flip came too early, another flip will be scheduled. If it came after the opponent's flip, the model records the estimated time between its own last two flips as an episode in declarative memory. It then uses a risk-adjusted blend of the previous period estimates to determine the time of the next flip. The flip is carried out as soon as its time has come (according to the temporal module).

Risk adjustment Both models use a risk adjustment procedure to bias an estimate μ, obtained in cognitive modality m, according to the function

$$\text{RiskAdj}(\mu, d, m) = \mu + d\frac{R}{E}\sigma_{m,\mu}$$

where $d \in \{-1, 1\}$ stands for the direction of risk (in our task the value is -1 as it is more risky to flip earlier). The model estimates the error associated with a modality m (visual, temporal) at an estimated size μ as $\sigma_{m,\mu}$. This knowledge is obtained via a long-term metacognitive process that monitors the reliability of one's own perceptual and cognitive processes.[3] $\sigma_{m,\mu}$ specifies the standard deviation of a gaussian distribution, indicating the error that is estimated for the specific cognitive modality. In accordance with the literature, we estimate it to be $0.015t$ for timing (t is the time in seconds), and $0.1s$ for spatial distances (s is the distance).

The acquisition procedure for R has the model make risky or conservative choices based on prior experience; outcomes are sampled from a normal distribution. Positive outcomes reinforce the chosen strategy. Early choices, risky or not, are hardened to form the model's long-term biased habits, as long as they lead to a few successful outcomes early on. The result is a distribution of risk propensity comparable to that observed in our data.

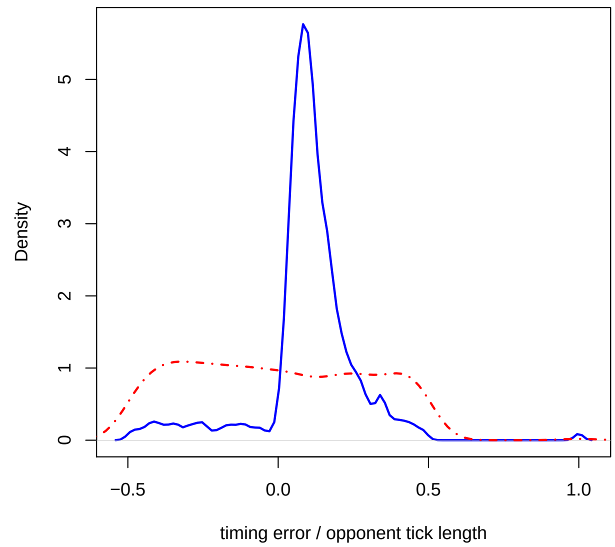

Figure 6: **Model** timing bias: as in Fig. 3, (**lower**-than-mean risk propensity personalities).

Discussion Figures 6 and 7 compare low and high risk propensity individuals, as simulated by the model. The graphs show a flat distribution for Exp. 1 for high-risk individuals between $[-.5, .5]$, but not for low-risk participants, and for Exp. 2 they show good risk management (a spike after the target timing, i.e., the opponent's flip). This reproduces our findings of the empirical data. Our model is exploratory and cannot provide a quantitative fit. In particular, we do not model the learning process across rounds. The empirical data reflect a complex 20-second task associated with deliberate reasoning and learning of the task at an early stage, a process that is rarely modeled in the literature. Regression models are made available in supplemental materials.

Conclusion

Risk is a function of probability and payoff. In a purely economic decision-making experiment, both are controlled and known to a subject. FlipIt introduces a cognitive variant of such experiments: while the payoff is externally controlled, probability is defined by the subject's internal estimation error. For this reason, the task seems more ecologically valid than many simple, serious games. With its short, 20-second rounds, it focuses on short-term learning. The core contributions we make are as follows.

Empirically, we show how risk propensity is related to a complex task that involves risk taking on an implicit basis.

[3]This predicts that a younger, less experienced individual will behave more erratic in choosing which risks to take, as σ has not yet converged.

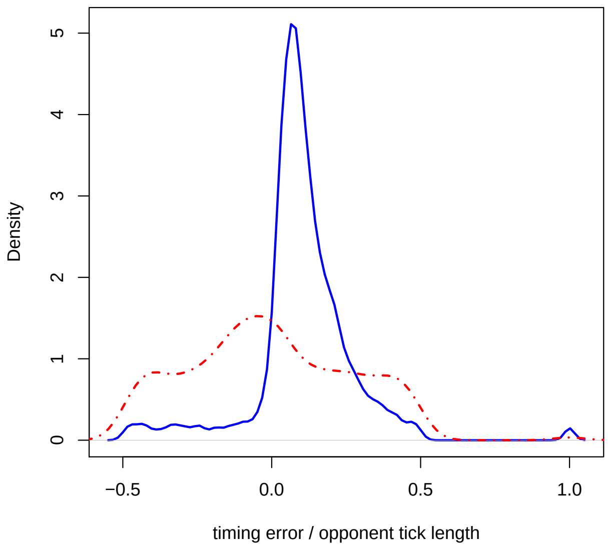

Figure 7: **Model** timing bias as in Fig. 4, (**higher**-than-mean risk propensity personalities).

Implicit, short-term processes do not adopt irregular underlying payoff or risk distributions from experimental descriptions. Instead, subjects apply a Gaussian heuristic, even if it contradicts the payoff structure of the game.

Our data demonstrate that high-risk individuals use risky behavior as a default. However, they eventually learn to manage the risk with respect to the specific task. Unlike experiments informing most cognitive models, we have run two designs that vary cognitive factors (internal vs. external memory, estimation of timing vs. estimation of spatial difference). The experiments point out dependencies in human decision-making that are uniquely suitable to modeling with a modular cognitive architecture.

Our cognitive model outlines possible mechanisms for these effects in reliance on established and validated estimation and perception error rates. It makes a developmental (rather than genetic) case for individual differences in risk propensity, and it shows how risk propensity may be combined with a metacognitive self-assessment of perceptual and temporal estimation error.

References

Blackwell, D. (1949). *The noisy duel, one bullet each, arbitrary accuracy* (Tech. Rep.). The RAND Corporation, D-442.

Brunnermeier, M., & Morgan, J. (2010). Clock games: Theory and experiments. *Games and Economic Behavior*, *68*(2), 532 - 550.

DellaVigna, S., & LiCalzi, M. (2001, January). Learning to make risk neutral choices in a symmetric world. *Mathematical Social Sciences*, *41*(1), 19-37.

Friedman, E., Shor, M., Shenker, S., & Sopher, B. (2004, May). An experiment on learning with limited information: Nonconvergence, experimentation cascades, and the advantage of being slow. *Games and Economic Behavior*, *47*(2), 325–352.

Gonzalez, C., Lerch, F., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591-635.

Kahan, J., & Rapoport, A. (1974). Decisions of timing in bipolarized conflict situations with complete information. *Acta Psychologica*, *38*, 183–203.

Kahneman, D., & Tversky, A. (1979, March). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263-291.

March, J. (1996, April). Learning to be risk averse. *Psychological Review*, *103*(2), 309-319.

Meertens, R., & Lion, R. (2008). Measuring an individual's tendency to take risks: The risk propensity scale. *Journal of Applied Social Psychology*, *38*(6), 1506–1520.

Meertens, R., & Lion, R. (2011). The effects of risk-taking tendency on risk choice and pre- and post-decisional information selection. *Journal of Risk Research*, *14*(6), 647-656.

Nochenson, A., & Grossklags, J. (2013). A Behavioral Investigation of the FlipIt Game. In *12th Workshop on the Economics of Information Security (WEIS).*

Radzik, T., & Orlowski, K. (1982). A mixed game of timing: Investigation of strategies. *Zastosowania Matematyki*, *17*(3), 409–430.

Rapoport, A., & Murphy, R. (2012). Evolution and breakdown of trust in continuous time. *The Oxford Handbook of Economic Conflict Resolution*.

Reitter, D., & Lebiere, C. (2010). Accountable modeling in ACT-UP, a scalable, rapid-prototyping ACT-R implementation. In *Proceedings of the 10th International Conference on Cognitive Modeling (ICCM)* (p. 199-204).

Salvucci, D. (2001). An integrated model of eye movements and visual encoding. *Cognitive Systems Research*, *1*, 201-220.

Taatgen, N., & van Rijn, H. (2011). Traces of times past: Representations of temporal intervals in memory. *Memory & Cognition*, *39*(8), 1546–1560.

Taatgen, N., van Rijn, H., & Anderson, J. (2007, July). An integrated theory of prospective time interval estimation: The role of cognition, attention and learning. *Psychological Review*, *114*(3), 577-598.

van Dijk, M., Juels, A., Oprea, A., & Rivest, R. (2012). *Flipit: The game of "stealthy takeover"* (Tech. Rep.). Cryptology ePrint Archive, Report 2012/103, 2012.

Weber, E., & Johnson, E. (2009). Chapter 10 - Decisions under uncertainty: Psychological, economic, and neuroeconomic explanations of risk preference. In P. Glimcher, C. Camerer, E. Fehr, & R. Poldrack (Eds.), *Neuroeconomics* (p. 127 - 144). London: Academic Press.

Wood, S., & Swait, J. (2002). Psychological indicators of innovation adoption: Cross-classification based need for cognition and need for change. *Journal of Consumer Psychology*, *12*(1), 1–13.

Zuckerman, M., Kolin, E., Price, L., & Zoob, I. (1964, December). Development of a sensation-seeking scale. *Journal of Consulting Psychology*, *28*(6), 477-482.

Supplementary Material

	E1 Emp.	E2 Emp.	E1 Mod.	E2 Mod.
(Intercept)	11.6868^{***}	11.1875^{***}	14.1374^{***}	17.6540^{***}
	(0.3787)	(0.4060)	(0.0751)	(0.0632)
info.treatment.id	-0.0015	0.0611		
	(0.1217)	(0.1281)		
(anchor)	0.0057^{***}	0.0019^{**}	0.0097^{***}	0.0065^{***}
	(0.0004)	(0.0006)	(0.0003)	(0.0002)
(tick)	0.0139^{***}	0.0154^{***}	0.0144^{***}	0.0259^{***}
	(0.0004)	(0.0006)	(0.0003)	(0.0002)
risk.prop	-0.0032	-0.0115	0.2572^{**}	-0.0282
	(0.0174)	(0.0191)	(0.0912)	(0.0745)
(round.num)	0.5364^{***}	0.6715^{***}	0.0434^{*}	0.0038
	(0.0231)	(0.0383)	(0.0183)	(0.0104)
(education)	-0.2033	0.1385		
	(0.1419)	(0.1488)		
risk.prop:(round.num)	0.0145^{***}	0.0151^{***}	-0.0007	0.0049
	(0.0022)	(0.0036)	(0.0212)	(0.0121)
(round.num):(education)	0.0745^{***}	-0.1819^{***}		
	(0.0188)	(0.0287)		
Log Likelihood	-24226.0460	-11439.4555	-7018.8805	-10724.6877
Num. obs.	8856	4070	3427	6109
Num. groups: session.id	148	82	151	150

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$ by Markov-Chain Monte-Carlo sampling.

Table 1: Regression models on empirical and model data. Reponse: round payment in cents. Given the exploratory nature of our inquiry, we did not reduce our models to ones that only include significant effects. All predictors in parenthesis were centered around 0. "tick" indicates the opponent's flip period, "anchor" is the time of the initial opponent flip. Info.treatment.id treats information available to subjects (from less information: low values to more information). Education is the participants survey measure of education (lower numbers: lower highest degree).

	E1/E2 Emp.
(Intercept)	0.6991^{*}
	(0.3056)
(fliptime)	-0.0018^{***}
	(0.0005)
(risk.prop)	0.0089
	(0.0299)
expE2	0.8575
	(0.5490)
(risk.prop):expE2	0.0264
	(0.0525)
expE1:(round.num)	-0.0356
	(0.1762)
expE2:(round.num)	0.0907
	(0.2714)
(risk.prop):expE1:(round.num)	-0.0021
	(0.0172)
(risk.prop):expE2:(round.num)	0.0567^{*}
	(0.0261)
Log Likelihood	-61727.2155
Num. obs.	12926
Num. groups: session.id	230

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$ by Markov-Chain Monte-Carlo sampling.

Table 2: Regression model showing the influence of risk propensity over experience (round number) in the two different experiments. Response variable: timing bias (* 100) normalized by the period of the opponent's flips.