

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Impatience, Risk Propensity and Rationality in Timing Games

Permalink

<https://escholarship.org/uc/item/05w5q6vt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 36(36)

ISSN

1069-7977

Authors

Ghafurian, Moojan

Reitter, David

Publication Date

2014

Peer reviewed

Impatience, Risk Propensity and Rationality in Timing Games

Moojan Ghafurian and David Reitter

moojan@psu.edu reitter@psu.edu

The Pennsylvania State University

University Park, 16802, USA

Abstract

Games of timing reflect dynamic decision-making under uncertainty, as it takes place in many real-world situations, including health care, safety and security. Rather than making discrete decisions, participants choose one or more points in time that determine the outcome. We study individual's biases and characteristics in such games of timing. We examine risk propensity as a personal preference affecting timing decisions and document a bias, *impatience*. Experiment 1 analyzes people's strategy in timing games in relation to a rational model. Contrasting two cognitive models suggests that individuals apply risk propensity to the probability distributions underlying short games and when unfamiliar with the situation, but that, over time, impatience takes over as a linear adjustment. In Experiment 2, impatient participants risk their incentive payment in order to play early, even if they receive no advantage from doing so.

Keywords: Decision-making under uncertainty; Rationality; Patience; Risk propensity; Timing games

Introduction

Many real-world decisions concern the timing of one's actions. Good examples are common security or safety scenarios that are subject to cost/risk tradeoffs—at which point in time should guards patrol a dangerous area? While driving, how often we turn to look if a bicyclist has started to move while we are waiting to enter an intersection? Decision-makers need to judge risk under uncertainty and anticipate an opponent's actions. Unlike in much work on decision-making, the decision *when* to act becomes just as important as *how to act*. Games of timing generalize such scenarios and allow us to investigate the pertinent question whether well-known effects from classical, discrete decision-making apply.

In the study presented in this paper, we use games to analyze timing behavior and correlate it with both estimated and survey-based measures of risk propensity (a preference for risk-seeking over risk-avoidance). In the first version of the game, a virtual opponent chooses the timing of a covert move. The player is to anticipate the opponent's move and make his own move as shortly as possible afterward. Pre-empting the opponent is ineffective and costly to the player.

Generally, all games we use involve stochastic, time-based moves by an opponent, and one game features a payoff distribution that is not symmetric between early and late action. This is a characteristic of many real-life situations (e.g., in security, safety or personal health, where acting early incurs a small penalty, but acting late leads to large losses). This game is representative of a wide range of situations. Some of these situations would be calculated and encoded in policies (i.e., what is the best interval for checking fire safety devices in a public place? How often and at what age should people undergo medical screening for hints of a disease?) Yet, people

often fail to follow the prescribed policy, and many real-life situations require good judgment on the spot.

Previous work has suggested that risk propensity, an individual variable, influences people's decisions in similar games. Risk propensity introduces a bias that leads people to make decisions that are contrary to their interests. Yet, the actual implementation of risk propensity is unknown. For timing decisions, we propose a related heuristic, impatience, which causes people to act earlier than is advisable. The heuristic does not take expected utility into account.

Related Work

Timing decisions have seen relatively limited attention compared to discrete decisions. We know that time is perceived differently by different people. E.g., impulsive individuals over-estimate the duration of time intervals, specially when they are waiting for a beneficial outcome of their action (Wittmann & Paulus, 2008). This makes timed decision-making particularly challenging for some people.

Decision-maker's motivations and understanding of the risks greatly influence the outcome of timing decisions. A case in point is health screening, where people differ greatly in their compliance with policies that optimize risks and benefits (Rakowski, Fulton, & Feldman, 1993; Sonnenberg & Beck, 1993). Understanding human biases in such timing decisions is key to successfully convince people to act rationally in preventative care.

Recently, Reitter, Grossklags, and Nochenson (2013) used the *FlipIt* game of "stealth takeover" (van Dijk, Juels, Oprea, & Rivest, 2012) to relate risk propensity to decision-making under uncertainty in a dynamic, adversarial environment. In this game, players have to estimate the timing of multiple opponent moves and follow them (but not precede them) as closely as possible. Experiment 1 of this paper is a controlled equivalent to the very first move in *FlipIt*. We will consider rational solutions to the two games as the ones that optimize expected utility given currently available information. This contrasts with game-theoretic predictions, where individuals should, in repeated games, tend to choose equilibrium strategies (see Teraoka, 1983, for a duel-type timing game).

Experiment 1

Experimental games allow us to observe behavior in situations comparable to those mentioned above. We generally operationalize the timing situations as searching for an unknown event (happening in a specific period of time) when efforts for searching and latency in finding the event are both costly. During a round of l seconds duration, participants play




Figure 1: Feedback shown when the game is current and it is at point B on the bar. The latest the player has searched is at point A, and as a result, no information about the situation after point A is revealed.



Figure 2: Feedback shown at the end of the game. At point A the event has occurred, however the searching attempt has not happened until point B.

against an opponent that triggers a covert event at one point during the round. (The opponent plays exactly once and its move time is drawn from a uniform distribution.)

Players start out with $l * 100$ points. At any time, they can spend 100 points to check whether the opponent has played. The game ends after a successful check, otherwise it continues. Each second of latency in catching the opponent costs 100 points. So, for a 10-second game, with the opponent playing at 6.0s, and the player checking twice at seconds 4.0 and 6.5, the payoff is $1000 - 200 - 50 = 750$ points. This creates an incentive to neither check too early, nor wait too long.¹

In each experiment, participants completed a survey with four demographic questions, three basic integrity questions, and seven risk propensity assessing questions (Meertens & Lion, 2008). The risk propensity questionnaire includes Likert-scale questions such as “I really dislike not knowing what is going to happen”, or “I do not take risks with my health.”

Method

73 volunteers recruited on Amazon Mechanical Turk (33 female and 40 male, age mean 32y, [19,64]) participated in the experiment. Each played 25 rounds of the game. Participants were randomly assigned to a group that determined game length (5, 10, or 15 seconds). They were given detailed instructions with additional worked examples (see Appendix, Figure 6). Visual feedback of the game that was given to participants is shown in Figures 1 and 2.

Participants played one unpaid practice round, which was included in the results, followed by 24 paid rounds and were incentivized with a show-up fee of \$0.25, and a performance bonus of \$.004 for each 100 points they earned (no less than 0.00 per round).

¹Risk-taking in discrete decisions is affected by the polarity of incentives (Tversky & Kahneman, 1986). We study applicability of this bias to games of timing in a forthcoming paper.

Rational Model

We consider a cognitive model for playing the game—a rational model. This strategy always maximizes the expected utility under the known facts of the game: the opponent has not made a move so far, and how much time remains.

Assuming a pay rate of one per time unit, the iterative solution for the expected utility of a move at time t given a previous, unsuccessful move at time t_p , for a game of maximum duration d is:

$$U(t, t_p) = -MoveCost + \sum_{k=t_p}^t \frac{1}{d-t_p} (k-t_p + d-t) + \max[0, \frac{d-t}{d-t_p} \max_{t < m < d} U(m, t)]$$

This function iterates over possible opponent move times k up to the proposed check time t , whose probability is the inverse of the remaining time $d-t_p$. The payoff for these consists of the initial period until the opponent move, and the time period after our own move. For the remaining probability mass (i.e., in case the opponent has not played yet), rational future moves are assumed.

The rational strategies for games of different durations are shown in Table 1, along with their expected utilities.

We hypothesize that participants apply a personal preference for risk-taking to the rational decision. The rational player model is a first step for us to describe how exactly such a personal preference would be applied consistently. In this game, this rational player would seek to optimize the expected utility. Unlike in standard decision-making tasks designed to contrast risk-averse or risk-seeking behavior, there is only one optimal choice, and any risk-seeking or risk-averse behavior will reduce the participants payoff.

Strategy Analysis

For each round of the game played by participants, we find the expected utility of the sequence of moves they played and compare them to the maximum expected utility played by our rational model. This provides a meaningful measure of rational play (Figure 3). We also compare the first move of participants to the first move of the rational player (Figure 4). Playing much earlier than the rational player implies risk-seeking. The game might be over sooner, with a higher payoff, but the expected utility is lower. Thus, the incentive discourages such play; we interpret any bias of early play as risk-seeking or impatient behavior. The number of checks gives a similar picture. As the potential gain for each second of the game covers the cost of one check, checking more than once per second will certainly lead to a zero payoff.

Results

Given the sequence of participants moves, the expected utility for each round can be calculated. Figure 3 shows the distribution of outcomes for 5, 10, and 15-second games.




Figure 3: Expected utility of sequence of moves played by participants with different risk propensities. 95% confidence intervals (dotted lines), obtained by bootstrapping.




Figure 4: Time of the first move for participants with different risk propensities in three different periods of Experiment 1. 95% confidence intervals (dotted lines), obtained by bootstrapping.

Table 1: Rational Strategies

Duration	No. Checks	Timing of Checks	Expected U.
5sec	1	(2.50)	274
10sec	2	(4.00, 7.00)	670
15sec	3	(5.00, 9.00, 12.00)	1096

Risk propensity as measured by the survey predicts the participants' strategies. Risk propensity was reliably associated with less rational play (lower expected utility) in games of length 15s ($\beta = -7.3, t = -3.63$), while it did not affect people's outcome in shorter games of length 5s ($\beta = 0.79, t = .37$) and length 10s ($\beta = -3.0, t = -1.39$). Participants reliably make their first move earlier than the rational model in games of all lengths (5s: $\beta = 0.7s, t = -3.21, CI = \pm 0.43$; 10s: $\beta = -1.0s, t = -4, CI = \pm 0.42$; 15s: $\beta = 1.61s, t = -7.78, CI = 0.40$). Overall, game duration is a reliable predictor of timing difference ΔT to the rational model (Intercept 0.206s, $\beta_{\Delta T} = -0.09s, t = -3.01, CI = \pm 0.059$).

Once game duration is taken into account, risk propensity reliably predicts timing relative to optimal timing ($\beta = -0.03s, t = -2.67, CI = \pm 0.023$).²

²For uniformity of analysis and presentation, all analyses are linear mixed-effects regressions, with random intercept by subject. Correlation between all predictors was below 0.1. Significance at $p < 0.05$. All confidence intervals 95% based on t ; in graphs: bootstrapped.

Risk-taking or Impatience?

The implied assumption so far has been that the gameplay is adversely influenced by risk-seeking behavior. In their play, participants exhibit risk-seeking even in the face of lower utility (i.e., their play becomes irrational). This strategy, if it can be called that way, would go beyond the contrast of risk-seeking and risk-avoiding behavior known from classical, discrete tasks (Iowa Gambling Task: Bechara, Damasio, Damasio, & Anderson, 1994), where the game is designed to keep expected utility constant across the strategy space.

Possibly, players use a heuristic to implement their risk-seeking preferences. Alternatively, another bias could be at play, such as *impatience*. The rationale for this bias could be to improve one's information basis, which has been proposed as a rational explanation of behavior in the Wason Card Selection Task (Chater & Oaksford, 1999; Wason, 1966).

In the remainder of this paper, we will contrast the two biases using linear regression, before presenting further empirical evidence to determine whether impatience is indeed a bias that impacts rational play.

We fit four linear regression models that each predict the timing of play. Model *A* applies risk propensity in relation to the calculated hazard rate. Model *B* applies impatience, playing earlier by a consistent amount of time. Model *C* is a baseline model that plays a fixed strategy as given in Table 1, regardless of individual personality traits. Model *D* combines both *A* and *B* models.

Models *A* and *B* each estimate a subject-specific variable

that influences the prediction by playing earlier. Model *B* implements a heuristic. It adjusts timing linearly and the adjustment is the same for all moves in all rounds played by a given participant. By contrast, Model *A* simulates a personality trait that adjusts a probability threshold. *A* plays as soon as the cumulative probability of the opponent making his move has reached this threshold. *A* updates its hazard rate estimate: during a round, the cumulative probability increases linearly with time. The hazard rate is $1/t_{rem}$, with t_{rem} indicating the time remaining since the last unsuccessful check. (We know that the opponent will play for sure, and we know that the opponent has not played up until our last check.) In other words, the player’s check adds information that allows him to increase the hazard rate estimate and set the cumulative probability to 0 at check time. Thus, model *A* adjusts its situation model as soon as new information is obtained. Individual risk propensity is applied to the hazard rate estimate. For comparability, even model *A* estimates the subject-specific risk propensity rather than using the risk propensity we have obtained from the survey.

Model specification ΔT is the difference between a subject’s move and the previous move (or 0 for initial move). All models predict ΔT . Models *A*, *B*, and *D* use a linear function of game duration, game round, survey-assessed gender, years of education and age. Model *A* also includes an interaction of remaining game length and the subject-specific, fitted risk-propensity variable. Model *B*, instead, includes the remaining game length and the subject-specific, fitted impatience variable (both as individual, linear predictors). Model *D* combines the variables and interactions from *A* and *B*. Model *C* always predicts the time to the next check according to the rational strategy for a game of the given length (Table 1).

Modeling results Impatience and risk propensity both make significant contributions to explaining the variance in timing decisions, across all conditions. We see that model *D*, which uses both risk propensity and impatience is the best model. Comparing shorter and longer games, the relative importance of risk propensity and impatience reverses. Perhaps unsurprisingly, when given more time, participant’s impatience becomes more indicative of behavior. However, given just five seconds they are able to consistently apply risk propensity to their decision-making. This becomes clearer when we consider learning over time. In early rounds (rounds 1 – 8, labeled “novice”), the risk-taking model outperforms the impatience model; this is reversed in late rounds (17 – 24, labeled “expert”), where, (according to AIC) the impatience model is $e^{14/2}$ times more likely than the risk-taking one. (Note that impatience being a better predictor for late rounds does not imply that subjects become more or less patient.)

Experiment 2

The first experiment demonstrates that impatience outperforms risk propensity as behavioral predictor in a timing task.

Table 2: Model fit criteria (modified Aikake Information Criterion, following Burnham and Anderson (2002)). Lower AICc values indicate a preferable model.

Model	AICc	Δ AICc	Log-Lik
FULL DATASET			
D: combined	45,970	0	-22,830
A: risk-taking	46,229	259	-23,035
B: patience	46,688	718	-23,268
C: baseline	49,341	3371	-24,668
5-SECOND GAMES			
A: risk-taking	10,274	0	-5,108
B: impatience	10,390	116	-5,170
C: baseline	11,288	1,014	-5,642
15-SECOND GAMES			
B: patience	20,910	0	-10,424
A: risk-taking	20,936	25	-10,433
C: baseline	22,433	1,523	-11,214
NOVICE			
A: risk-taking	19,577	0	-9,704
B: impatience	19,672	96	-9,755
C: baseline	21,070	1,493	-10,532
INTERMEDIATE			
A: risk-taking	12,634	0	-6,255
B: impatience	12,660	26	-6,273
C: baseline	13,572	938	-6,784
EXPERT			
B: impatience	13,440	0	-6,663
A: risk-taking	13,454	14	-6,666
C: baseline	14,547	1,107	-7,272

In that task, participants had to understand the payoff distribution and the fact that the opponent will make a move for sure. This is crucial for our analysis. Because the opponent will play before the end of the game, the duration until the end of the game influences the probability density (i.e., probability that the opponent will play in the next second). As a consequence of constant individual risk propensity, participants should adjust their moves. Because the payoff dropped off linearly after the opponent’s move, participants were, to some extent, incentivized to make early moves. This ecologically valid element of the game leaves us with a question—did participants play too early because of their interpretation of the rules of the game?

Experiment 2 is designed to measure impatience in timing games more directly. In this task, we describe a similar situation in which participants are to intercept an external, stochastic opponent. However, they are either to ensure that they play after the opponent (“late” condition) or that they play before the opponent (“early”). Participants and opponents play only once. They are not incentivized to play as close as possible to the opponent, which makes the rational solution of the game trivial—play either very late, or very early, in order to reap the

full bonus. An impatient participant will always tend to play earlier rather than later. Thus, an impatient participant tends to do better in the “early” condition compared to the “late” condition. Participants were given the same set of survey questions as in Experiment 1.

Method

123 volunteers recruited on Amazon Mechanical Turk (43 female, 79 male, 1 unknown; mean age 31y, [18;67]) participated in the experiment for payment. Each played 10 rounds of the game. Participants were randomly assigned to a group that determined game length (5 or 15 seconds). In a between-subjects design, participants were assigned either the “early” or “late” condition. Visual feedback of the game was similar to that in Experiment 1. The procedure was the same as the previous experiment, except for the number of rounds and not having a practice round.

Unlike in Experiment 1, participants only play once per round, at which point any previous opponent move is revealed. Participants then see an animated progress bar until the end of the round; any future opponent move is revealed at its scheduled time. (Time is money, particularly on Mechanical Turk, and we did not want to incentivize playing early.)

Participants were paid a show-up fee and a bonus. The bonus was either the full amount for a win, or nothing otherwise. A win required them to play before the opponent in the early game, or after the opponent in the late game. The actual bonus payoff depended on the length of the game (100 points per second). (See Figure 7 for the scenario.)

Results

We contrast the time delay between beginning of the game and play in the early condition, and the delay between play and end of the game in the late condition, ΔT . Rounds in which participants did not play at all were excluded. Participants were better at minimizing the initial delay in the *early* condition than they were at minimizing the final gap in the *late* condition. This means, they showed a consistent preference to play early. Figure 5 illustrates this difference over the 10 rounds of the game, in relation to the length of each game (5 or 15 seconds). The difference was reliable for rounds 1 through 6 (Wilcox, $W \geq 1999$, $p < 0.01$).

Discussion

It is possible that a risk mitigation strategy supported the observed effect. Participants could play as soon as the game started in the *early* condition, but they had to avoid missing the end of the game in the *late* condition in order to score points. However, we point out that 1) participants could theoretically play too early (before the game started) and would not receive visual confirmation, and that 2) the performance gap between the two conditions is too large to be explained by such a strategy.

Is impatience a heuristic for risk-seeking? Based on the model comparisons for Experiment 1, we can state that both




Figure 5: In both conditions, early and late, participants showed a learning effect, and a persistent tendency to play too early rather than to late. 95% confidence intervals obtained by bootstrapping.

impatience and risk propensity make significant contributions to explaining the variable in the decisions made. However, are there commonalities? As a second test, we would hypothesize that risk propensity, as measured by the standard survey, would predict timing in Experiment 2, our impatience task. This is not the case—overall correlation of risk propensity and ΔT is very low ($\rho = -0.02$). Split by condition and game duration, the associations are also low ($|\rho| < 0.076$).

General Discussion

Our standard to compare to is rationality. Participants must time their moves in Experiment 1 according to the cumulative probability of the opponent’s move. While participants seem to approximate this model, we have begun ruling out simpler heuristics—impatience in our study—that might account for the observed variance. Future work will distinguish the rational model from other shortcuts. Another prediction of the rational model would be that participants will behave differently depending on whether the game has a known end point and the opponent will actually move. With an infinite game, or the possibility that the opponent will not move at all, a shorter remaining time span would not warrant shorter move intervals. If the *Gambler’s Fallacy* (Tversky & Kahneman, 1974) applies to timing decisions, we would expect participants to not change their behavior. Risk propensity, in this case, would still consistently act on the assumed, wrong probability density of the opponent’s move.

We see impatience as a personal bias that does not de-

pend on underlying probabilities, but may be influenced by the payoff and framing. While we set out to show that impatience is a heuristic that implements risk-taking in the kinds of situations reflected in timing games, our results led us to the opposite conclusion. Risk-taking and impatience affect decision-making independently.

The contributions of risk-taking and impatience to explaining the timing behavior change and depend on the length of the game and participants' progress in the game. Depletion of self-control (Baumeister, Vohs, & Tice, 2007) might play a role in making impatience a better predictor in later rounds of the game, assuming it applies to very short games like ours. However, as participants are performing better in the latter rounds (figure 5), we find that unlikely.

Conclusion

Timing decisions are of critical importance in many domains of decision-making; they are under-studied and pose a new avenue for understanding personality characteristics and general biases. Two related examples of these properties were studied using timing games. Risk propensity is an individual variable that can be estimated or obtained via a standard survey instrument: in both cases it predicts a willingness to take one's chances. People appear to make rational decisions over a broad range of risk propensity levels found in the population. Given enough time, the most risk propense people appear to optimistically bet on a positive outcome, resulting in less rational choices.

Most irrationality in the timing decisions we study appears to stem from a second characteristic: impatience. Acting as a bias regardless of the underlying probability distribution of winning, impatient subjects were willing to forego certain winings in exchange for seeing the outcome of the game earlier, or simply for locking in their decision.

A precise picture of how personal traits and common biases influence timing decisions will be of use in designing training problems for a range of people whose successful, rational decision-making can make a difference in domains such as personal health, national security, or public safety.

Acknowledgements

We thank Jens Grossklags and Frank Ritter for comments and Ethan Heilman and Alan Nochenson for their FlipIt code.

References

Baumeister, R. F., Vohs, K. D., & Tice, D. M. (2007). The strength model of self-control. *Current directions in psychological science*, 16(6), 351–355.

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1), 7–15.

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multi-model inference: a practical information-theoretic approach*. Springer.

Chater, N., & Oaksford, M. (1999). The probability heuristics model of syllogistic reasoning. *Cognitive Psychology*, 38(2), 191–258.

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

Rakowski, W., Fulton, J. P., & Feldman, J. P. (1993). Women's decision making about mammography: a replication of the relationship between stages of adoption and decisional balance. *Health Psychology*, 12(3), 209.

Reitter, D., Grossklags, J., & Nochenson, A. (2013). Risk-seeking in a continuous game of timing. In *Proceedings of the 13th International Conference on Cognitive Modeling (ICCM)* (pp. 397–403).

Sonnenberg, F. A., & Beck, J. R. (1993). Markov models in medical decision making a practical guide. *Medical Decision Making*, 13(4), 322–338.

Teraoka, Y. (1983). A two-person game of timing with random termination. *Journal of Optimization Theory and Applications*, 40(3), 379–396.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, S251–S278.

van Dijk, M., Juels, A., Oprea, A., & Rivest, R. L. (2012). FlipIt: The game of “stealthy takeover”. *Journal of Cryptology*, 1–59.

Wason, P. C. (1966). Reasoning. *New Horizons in Psychology*, 1, 135–151.

Wittmann, M., & Paulus, M. P. (2008). Decision making, impulsivity and time perception. *Trends in Cognitive Sciences*, 12(1), 7–12.

Materials

You have invited cookie monster over. Unfortunately, you have a good number of cookies in your living room, where he is waiting while you cook dinner in the kitchen. Cookie monster will start eating the cookies at one point, so you need to check on him every now and then. The sooner you catch him, the more cookies you save, however, anytime you check on him and your friend has not started eating, he'll get more and more annoyed with you. You will need to give him a pack of 100 cookies every time you check on him.

Figure 6: Exp. 1, Cover Story (abbreviated)

Imagine you are Cookie Monster! You really love cookies. Your friendly neighbor knows that well, so she has promised to bring you some cookies. The problem is that she hasn't told you when she will bring the cookies. It's a very cold day, and you can only open your door once and check very quickly if the cookies are there. Your job is to figure out when to check in front of your door. Remember: Your neighbor will definitely put the cookies for you there within the promised period, the question is: when? ... If you don't find the cookies, you won't get them at all. If you check before your neighbor brings the cookies, you have lost your chance to get them.

Figure 7: Exp. 2, Cover Story (“late” condition, abbreviated)