



Learning through period and physical time[☆]

Teresa Backhaus^a, Steffen Huck^{b,c,*}, Johannes Leutgeb^b, Ryan Oprea^d

^a University of Bonn, Institute for Applied Microeconomics, Adenauerallee 24–42, 53113, Bonn, Germany

^b Wissenschaftszentrum Berlin für Sozialforschung, Reichpietschufer 50, 10785, Berlin, Germany

^c University College London, Department of Economics, Gower Street, WC1E 6BT, London, United Kingdom

^d University of California Santa Barbara, Department of Economics, 3014 North Hall, 93106–9210, Santa Barbara, CA, United States



ARTICLE INFO

Article history:

Received 27 October 2022

Available online 7 June 2023

JEL classification:

C73

C92

Keywords:

Cournot oligopoly

Laboratory experiment

Learning

Time

ABSTRACT

We demonstrate in a laboratory experiment in which subjects play a two-player Cournot–Tullock game over hundreds of periods of varying length that full accounts of subjects' learning requires the consideration of, both, 'period time' and 'physical time.'

© 2023 Elsevier Inc. All rights reserved.

1. Introduction

We document, using a laboratory experiment, that strategic learning with high-frequency interaction depends on two aspects of time – period time (the number of stage games that have elapsed) and physical time (the conventional “clock” time that has elapsed). The extensive literature on strategic learning has so far only considered period time, modeling agents as learning from interactions in past periods to make decisions in the current period using a variety of learning rules.¹ We show that this approach to time may be sufficient for understanding learning in some circumstances but that in others a full account of subjects' learning requires the additional consideration of physical time.

Our data suggest that as long as subjects have access to an appealing learning rule that yields satisfactory and stable payoffs, they see little reason to revise their behavioral rule. As a consequence, behavioral dynamics can be well captured by exclusively studying how behavior evolves over periods. If, in contrast, subjects are drawn to initial learning rules that produce poor or unreliable payoffs they start to question their initial rule and begin to experiment with other modes of behavior. This process of rule-adaptation requires careful reasoning which, in turn, renders physical time important in

[☆] Financial support by the Deutsche Forschungsgemeinschaft (CRC TRR 190 and CRC TR 224) is gratefully acknowledged.

* Corresponding author at: Wissenschaftszentrum Berlin für Sozialforschung, Reichpietschufer 50, 10785, Berlin, Germany.

E-mail address: steffen.huck@wzb.eu (S. Huck).

¹ Learning in period time has essentially been modeled as a mechanical process: inputs from the past generate outputs in the present through the use of a heuristic. Such heuristics include myopic best replies (Cournot, 1838), fictitious play (Brown, 1951; Fudenberg and Kreps, 1993; Boylan and El-Gamal, 1993) and other types of forecasting (Bao et al., 2013), satisficing behaviour (Simon, 1972; Nowak and Sigmund, 1993), conformity (Asch, 1952; Huck et al., 2002), reinforcement learning (Erev and Roth, 1998; Camerer and Hua Ho, 1999) or imitation (Vega-Redondo, 1997; Apesteguia et al., 2007).

shaping behavioral dynamics. Thus, in settings in which not only behavior but the rule driving behavior changes over time, period time alone is insufficient for understanding learning.

Our results have important consequences for both the theory of learning and for experimental methodology. If learning is confined to the adaption of choices from one period to the next, with a fixed mode of behavior, existing theories of learning are perfectly adequate to capture the data and experimenters can conveniently study longer time horizons within the same span of physical time by making periods shorter and shorter – within bounds of reasonable reaction time. If, in contrast, subjects also engage in the adaption of their adjustment rules, theories of learning must also account for the passing of physical time and experimenters can no longer simply ‘speed things up’ in experiments by shortening periods (or by studying continuous time settings, e.g. Friedman and Oprea (2012)).

In order to identify the role of physical time in strategic learning, we study the same underlying game under two information conditions and employ three different period lengths. Specifically, we implement an experimental 2-player Cournot-Tullock game that is characterized by a tension between competition and cooperation. The previous literature on such games has documented the important role of information conditions for subjects’ behavior. In particular, it has been shown that the learning rule subjects use depends on whether subjects have ready access to information about their payoff function that allows them to employ variations of best-reply rules to adapt their choices. With such information subjects tend to rely on myopic best-reply rules and behavior rapidly converges to the Cournot-Nash equilibrium. This has been shown for conventional experiments where subjects have arbitrary amounts of time to make their decisions and that last for less than 100 periods (Huck et al., 1999; Offerman et al., 2002) and also, more recently in experiments with high-frequency interaction where period lengths are given and short and that last for over 1,000 periods (Huck et al., 2017).

Importantly for our purposes, subjects use a very different behavioral rule when they do not have access to their payoff rule and are forced to rely instead on other subjects’ past actions and payoffs. In such low-information environments subjects tend initially to imitate successful actions taken by other players which pushes them, as predicted by theory (Vega-Redondo, 1997; Apesteguia et al., 2007), to very competitive outcomes and eventually very low payoffs. In conventional experiments with a relatively small number of periods subjects never drop their imitative behavior and outcomes remain highly competitive until the very end of session. However, when play is sped up in high-frequency experiments, where subjects make more than a thousand decisions in rapid succession, subjects eventually drop the imitation rule and switch to different modes of behavior that align the players’ incentives by matching quantities, using a form of punishment and forgiveness which, in the long run, lead them to cooperative behavior, as first shown in Friedman et al. (2015).

We use the known contrast in behavior across these two settings to pose our motivating question. Specifically, we compare subjects’ behavior in these two informational environments in a high-frequency setting and independently vary the period length (physical time, between two, four, and eight seconds).

On the one hand, we conjecture that period length will be largely irrelevant when subjects have easy access to payoff information in the high-information environment and can use a stable best reply learning rule. Since previous studies do not find changes in the behavioral rule over time, regardless of whether subjects play under a 100 or more than 1,000 periods, we speculate that subjects will not need to use physical time to re-formulate their behavioral rule. On the other hand, we conjecture that in the low-information environment where subjects tend to initially choose a poor (and in the long run unsustainable) behavioral rule, physical time may be important in shaping learning.

These conjectures are born out by the data. While period length does not play a significant role for dynamics when subjects are given sufficient information to best reply, we do observe significant differences between treatments with differing period lengths in the low-information environment, thus, establishing, to the best of our knowledge for the first time, the relevance of physical time for strategic learning.²

2. The experimental environment

We study 2-player Cournot-Tullock games that are characterized by a strong tension between cooperation and competition, with Nash equilibrium somewhere in the middle between full cooperation (the symmetric joint profit maximum) and full competition (the Walrasian outcome where profits are minimized). The payoff function we implement is

$$\pi(a_i, a_{-i}) = 10 + \left(\frac{120}{\sum_j a_j} - 10 \right) a_i$$

² The past literature on experimental games has only provided some indirect evidence for the role of physical time for learning. For example, Rick and Weber (2010) show that subjects acquire notions of dominance reasoning in environments with multiple periods without feedback. As subjects do not receive inputs from the past learning in their experiment must, hence be due to a process of thinking taking place in physical time. Similar effects have been documented in experiments where subjects make decisions in groups and are given the chance to deliberate (see, for example Goeree and Yariv, 2011). Other studies in experimental economics that examine aspects of physical time include the large body of investigations into temporal discounting surveyed by Frederick et al. (2002) and ever growing since examinations of the role of reaction time (Rubinstein, 2007; Rand et al., 2012; Krajbich et al., 2015; Spiliopoulos and Ortmann, 2018) and, more recently, perceptions of time (Brocas et al., 2018). None of these studies speak, however, to strategic learning in games.




Fig. 1. The Interface in the HIGH treatments. The interface in the LOW treatments is identical except that we omit the black line showing subjects their payoffs at every possible quantity choice they could have made given their partner's action in the previous period. This additional piece of information makes the myopic best response salient. On the very top subjects see how many periods are remaining in the game, how many points they have earned so far, and a progress bar showing when a period ends. Subjects can make choices using the rectangle at the bottom of the graph. Subjects receive feedback on their actions and earnings (green mark and dot respectively) and on their partner's actions and earnings (red mark and dot respectively) in the previous period. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

where $a_i \in [0.1, 6]$ is the own quantity choice and a_{-i} the opponent's quantity choice. The Nash equilibrium is at 3, the fully competitive (Walrasian) outcome is at 6 and the symmetric joint profit maximum at 0.1. The corresponding market profits are 80, 20 and 138 respectively.

We explore three different time treatments, each lasting for a total of 4,800 seconds of physical time. The treatments differ in their period length, λ , with $\lambda \in \{2, 4, 8\}$ seconds, which result in high-frequency games with 2400, 1200, and 600 periods, respectively. We run these treatments in two different information environments that differ in the amount of information that subjects have when making a decision. In both environments subjects see on their screens the actions and resulting payoffs for both players from the previous period. In treatments labeled LOW- λ , (with λ indicating period length) this is the only feedback information that subjects receive. In treatments labeled HIGH- λ subjects observe, in addition, a curve that traces all possible payoffs they could have been achieved in the last period as a function of each possible action the subject might have made, conditional on their counterpart's actual action.³ In the first round the computer assigns subjects a random initial strategy that subjects are forced to play. Afterwards, subjects are free to choose any action from the interval $[0.1, 6]$.

Fig. 1 shows the interface that subjects use. In any given period subjects make their choice by moving the rectangle at the bottom of the screen to a position on the x-axis. As long as the period is running they can move the rectangle freely to any position they want but the action is not implemented until the period ends (times out). Furthermore, they receive no feedback on their partner's action until the period is over. On the top of the interface a bar fills up showing the passage of time during the period: as soon as it fills up, the subjects' actions get implemented and a new period starts seamlessly. Subjects' choices are carried over from one period to the next unless changed. When the period ends, subjects receive feedback about the results in the previous period. They see a green mark on the x-axis showing their own action and a red mark showing their partner's action in the previous period. In the graph, above the marks, a green dot shows their earnings and a red dot shows their partner's earnings in the previous period. Furthermore, in the HIGH treatments, there is a black line showing subjects what they could have earned at each choice given their partner's choice in the previous period.

³ Huck et al. (2017) examine the short- and long-run effects of the different information environments and show that additional information about counterfactual payoffs inhibits cooperation in the long run. They do not investigate the role of period vs physical time though. Providing more illustrative instructions may affect behavior in itself as subjects with more detailed instructions may get better results Charness and Levin (2005).

Experiments were conducted at the WZB-TU laboratory in Berlin and were implemented using the ConG software package (Pettit et al., 2014), which allows subjects in the experiment to interact for many short periods in rapid succession. Subjects were recruited from the WZB-TU subject pool via ORSEE (Greiner, 2015). For each of the treatments we observe 36 subjects for a total of 216 subjects. In each session there were 12 subjects. After entering the laboratory, subjects were randomly seated at a computer and received written instructions. The instructions introduce the subjects to the graphical interface and, specifically, how they make their choices and what feedback to expect. Subjects know that they are matched with another participant for the entire duration of the experiment. Subjects are explicitly told that the payoff function mapping their choices into payoffs is hidden from them. The instructions stress that payoffs depend on the actions of the participant and the participant's partner only, that the payoff function stays the same for the entire experiment and that the payoff function is symmetric for both players. The entire instructions are reported in the Appendix. Since in treatments with higher λ subjects play more repetitions of the game, subjects' points were divided by $\frac{\lambda}{16000}$ to keep incentives constant over time. All sessions lasted 80 minutes (or 4800 seconds) and subjects earned on average EUR 14.52 plus a show-up fee of EUR 5. All earnings were paid out in cash at the end of the experiment.

3. Results

As discussed above, we conjectured that physical time (i.e. period length) would not play an important role in HIGH but may have an impact on dynamics in LOW. The conjecture is based on the observation that subjects in previous experiments tend to use relatively stable myopic best reply rules in environments like HIGH and therefore need little physical time to reformulate their decision rule. By contrast subjects in environments like LOW tend to change their decision rule with experience, abandoning the imitation rule subjects tend to initially use. The amount of physical time they have to contemplate and reformulate their decision rule may therefore influence the rule they replace imitation with.

In order to investigate this, we compare the evolution of market quantities over physical time (seconds) and over period time (periods) separately. Both perspectives on the data are shown in separate panels in Fig. 2. When we examine periods (Panel 2a) we focus on the first 600 (to match the treatment with the fewest number of periods) and divide the data into four 150-period segments. When we focus on physical time (Panel 2b) we instead divide the data into four brackets of 1200 seconds each. We focus on our 2- and 8- second conditions where our comparisons are best powered (4-second data is generally intermediate and is reported in the Appendix). For each time bin we plot a bar for the average in the 2-second and 8-second conditions, each with bootstrapped 95% confidence intervals.

In order to interpret these figures we note (i) that if only period time matters (as is generally assumed in the theoretical literature) there should be no significant differences between the 2- and 8-second treatments in Panel 2a, which organizes by periods, and (ii) that if only physical time matters for the evolution of play there should be no differences in Panel 2b, which organizes by seconds.

We begin with treatment HIGH (right panels) where we expect that physical time will not play a major role (we discuss the LOW treatment separately, below). We see, as in past experiments, that quantities slowly fall towards the Nash equilibrium over time in HIGH. Regarding (i) it is easy to see that there are indeed no meaningful differences between the HIGH treatments when we organize by periods (Panel 2a). Any pairwise comparison⁴ in any of the four bins between the 2-second and the 8-second treatments is statistically insignificant, indicating that the length of periods does not matter in HIGH.⁵ There are, however, differences when looking at the data through the lens of physical time. In any physical time window after the first bin, subjects who have experienced more periods (the 2-second condition) are further along in their procession towards the Nash equilibrium (Panel 2b).⁶

What does this mean? Because there are no meaningful differences between the two treatments across the period dimension in the HIGH environment, the data in HIGH can indeed be organized by ignoring physical time. However, inspection of Panel 2b suggests that the total number of periods that have passed does matter. Specifically, we see that, as the horizon gets longer, play does get more cooperative. For the theorist this implies that for an environment like HIGH there is no need for a fundamental revision of learning theories (as in most models, only period-time matters) and for the experimenter it means that she can indeed study longer period-time horizons by making periods shorter and, moreover, that these longer period-time horizons do matter for behavior. Longer period-horizons render play more cooperative and that increase in cooperation is entirely summarized by measures of period-time.

Next we turn to the LOW treatments. Here, a completely reversed picture of the role of time emerges. As in prior work, quantities fall slowly towards collusive levels over physical time. However, while there are no meaningful differences between the 2- and 8-second treatments as a function of physical time⁷ (Panel 2b), there are substantial and highly significant differences when we look at the data only through the lens of period time⁸ (Fig. 2a).

⁴ In this section we calculate the average market quantity for each group within each bin and test for differences between the treatments using Mann-Whitney-U tests. We report the p-values.

⁵ $p = 0.988$, $p = 0.323$, $p = 0.839$ and $p = 0.462$ for bins 1 through 4.

⁶ $p = 0.462$, $p = 0.091$, $p = 0.055$ and $p = 0.059$ for bins 1 through 4.

⁷ $p = 0.888$, $p = 0.389$, $p = 0.825$ and $p = 0.888$ for bins 1 through 4.

⁸ $p = 0.004$, $p = 0.016$, $p = 0.027$ and $p = 0.133$ for bins 1 through 4.




Fig. 2. Average Market Quantities. Panel a shows data for the first 600 periods and Panel b shows data for all 4800 seconds, split into four quarters each. For each quarter we calculate the average market quantities and bootstrap 95% confidence intervals.

The empirical pattern and therefore conclusions for LOW are exactly opposite to those we reached above for HIGH. It is not only that in low-information settings physical time matters for behavior – it appears that in LOW *only* physical time matters for behavior. That is, holding the number of seconds that have elapsed in the game so far constant, behavior does not seem to depend at all on the number of periods that have elapsed so far. This would seem to have important theoretical implications as it suggests that time matters for the evolution of behavior in ways not well-summarized by the number of repetitions and revelations of mutual behavior (i.e. periods) that have elapsed. Methodologically it produces an important caution: an experimenter cannot costlessly ‘speed up’ experiments to study the very long term by making periods shorter. Or alternatively, physical time is not a neutral consideration for experimental designs focused on low-information learning: when we expand or contract the length of periods in such settings, we significantly change behavior.

4. Conclusion

Our experiment is simple and our results are therefore easy to summarize. We experimentally study the long run play of a simple repeated social dilemma and vary (i) whether repetitions (periods) are short (2-second) vs. long (8-second). We show that in settings in which subjects are explicitly given direct information about the underlying payoff function (HIGH treatments), they make similar decisions (and follow similar long run behavioral trajectories) regardless of the length of the periods or how much time has elapsed in the game so far. However, when the underlying payoff function is not explicitly described (and must instead be inferred), we find exactly the reverse (LOW treatments). The number of periods that have elapsed is incidental in low information settings – rather the amount of physical time subjects have had to reason about the game so far is all that matters for describing the evolution of cooperation.

These results suggest that the theories of learning we usually employ when studying strategic interactions in economics are in an important sense incomplete. There are settings in which, indeed, learning depends only the elapse of periods (i.e.

the number of times agents' actions and resulting outcomes have been mutually revealed) as most theoretical models would suggest. But this seems to be a special case. In many (perhaps most) realistic settings information about the structure of payoffs is incomplete (sometimes radically so) and in such settings, physical time can matter as much as the conventional accounting of time employed in most strategic learning theories in economics.

Why should this be true? A natural interpretation is that implementing an already adopted learning rule (or mental model) is rather easy but formulating or discovering new rules and models is costly and takes time. In our HIGH treatment subjects have access to a natural initial learning rule – myopic best reply – suggested directly by the information we give them about their payoff function after each period. Following this rule, as subjects seem to, leads to progressively higher payoffs and so subjects are never inspired to reconsider that rule. By contrast, in our LOW treatment the most salient rule available is to “imitate the best” (that is, to match the quantity of whichever competitor earned the most in the previous period) – an intuitive rule that leads subjects to increase their quantities and lowers their payoffs over time. As in most low information Cournot experiments we find initial evidence of this, but subjects realize quickly that it is a bad rule and are forced to re-consider. In the absence of any other guides to decision making provided by knowledge of the mechanism, subjects are forced to consider and reconsider how they learn and adapt over time. While implementing a rule is largely automatic and independent of physical time, reformulating the bases of one's behavior requires contemplation and therefore instead outlays of physical time.

What are the implications of these findings? First, methodologically, our experiment raises some important lessons for the design and interpretation of experiments. Most importantly, although most theoretical models are neutral with respect to physical time, actual behavior is not. This can lead to serious distortions in the evaluation of behavior and estimation of decision rules if not carefully attended to. For example, the treatment difference between our HIGH and LOW treatment over the course of periods is radically larger in 8-second than 2-second periods even though, under the lens of standard learning theories, they are exactly the same. This means that using time-constrained experiments in order to study e.g. long-run learning must be used cautiously as they materially influence the degree of learning subjects accomplish.

Second, theoretically, our results highlight that learning has a multi-dimensional character rooted in both the pace with which information is revealed (period time) and the raw clock time required to perform cognitive acts. While automatic behaviors (i.e. implementations of heuristics and other rules) depend on the former, complementary cognitive acts required to formulate these rules depend on the latter. Developing new models of the interaction between these two types of learning, rooted in these two notions of time, may produce important insights not only on how mutual strategic behavior unfolds but also on how people learn from experience in individual efforts to optimize.

Declaration of competing interest

Authors declare that have no conflicts of interest.

Data availability

Data will be made available on request.

Appendix A. Methods

A.1. Experimental setup

The experiment was run employing the ConG software package (Pettit et al., 2014). In each of the treatments we observe 18 independent anonymous pairs of subjects who play the game for 4800 seconds. Subjects were recruited via ORSEE (Greiner, 2015) at the WZB-TU laboratory in Berlin. Subjects' average payouts were 14.52, plus a 5 Euro show-up fee. Upon arrival subjects were randomly allocated to seats and received written instructions explaining (1) they would play against one other participant; (2) how they would choose from their action spaces; and (3) what feedback they would receive through the graphical interface. Subjects are not made aware of the payoff function. They only know that the payoff function depends exclusively on both players' choices and that it is symmetric and does not change over time. The experimental instructions are reported below.

A.2. Code availability

The computer code used in the analysis is available from the corresponding author on request.

A.3. Data availability

The data that support the findings of this study are available from the corresponding author on request.




Fig. A.3. The Interface.

A.4. Detailed instructions

This section contains the instructions that subjects received in the HIGH treatment. The instructions are translated from German for review purposes. Instructions in the LOW treatment are the same except for a different screenshot and minus the sentence mentioning the black line.

Instructions Welcome! Thank you for participating in this economic experiment. If you read these instructions carefully, you can earn a non-trivial amount of money. The money that you earn during the course of this experiment will be paid to you in cash at the end of the last period. Please remain quiet and do not look at the screens of the other participants. If you have questions or if you need help, please give us a hand sign and we will come to your place. If you disrupt the experiment by speaking, laughing, et cetera, we will exclude you from the experiment without payment. We expect and appreciate your cooperation. All procedures in the experiment will take place exactly as they are described in these instructions.

Basic structure of the experiment In this experiment the computer will match you anonymously with another player. The experiment is divided into periods. In each period you and the other player will secretly choose actions. The combination of actions that you and your partner have chosen at the end of the period will determine the amount of points that you earn in this period. We will not explain to you exactly how your points are calculated, but here are some hints:

- Your points in each period are determined solely by your strategy and the strategy of your counterpart.
- The function that determines your points will not change during the experiment. If you and your counterpart choose the same actions at some points in time A and B, you will earn the same amount at point A as in point B.
- Your payoff function is symmetric to your counterpart's function. If you and your counterpart choose the same action in the same period, you will earn the same amount of points.

Computer display Fig. A.3 shows the display which you will use to make choices and through which you will interact with your counterpart. At the top of the display you see a progress bar that shows how much time has passed in the current




Fig. A.4. Average Market Quantities. Panel a shows data for the first 600 periods and panel b shows data for all 4800 seconds, split into four quarters each. For each quarter we calculate the average market quantities and bootstrap 95% confidence intervals.

period. When the bar is full the period ends and another period starts immediately. Your action is the position (from the left to the right) of the black square at the lower part of the display. During a period you can change your preliminary action freely by moving the square like a slider to the left and to the right, or by clicking on the desired position. Your actual action for the entire period is only determined by the position of the slider at the end of the period.

After a period has ended you will see a green point that shows the amount of points that you earned in the previous period. The higher the point the more points you have earned. The exact number of points is shown next to the point. At the same time you will see a red mark at the bottom of the display which shows the action of your counterpart in the previous period. You will also see a red point that shows the amount of points that your counterpart earned in the previous period. Next to this point you will also see the number of points that your counterpart earned. Finally you will see a black line that shows you how many points you could have earned at each position of the slider, depending on the action that your counterpart has chosen in the previous period. It is important for you to understand that the action of your counterpart, your points, and the points of your counterpart are always the results of the previous period. You will not receive any information about the points or the action of your counterpart in the current period, until it has passed.

Earnings In this experiment you will first earn points that are then converted into Euro at a rate of 0.3 Euro per point and paid out to you in cash. The exchange rate is noted down on the whiteboard at the end of the room. The earnings that are shown to you at the end of the period are the amount of points that you would earn in the entire experiment if you and your counterpart would decide the same in all periods. Your points will accumulate over the course of the experiment. The points that you have already earned are shown at the upper end of the display.

If you have not understood something, please raise your hand. We will answer your questions personally. Thank you for your participation!

A.5. Additional treatment

Fig. A.4 shows the same data as the figure in the main text but includes an intermediate treatment with 4-second period length. A consistent picture emerges as both panels show the 4-second treatment generally between the 2- and 8-second treatments, but the confidence intervals for the 4-second treatment typically overlap with the confidence intervals of the more extreme treatments due to lower power.

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.geb.2023.05.008>.

References

- Apesteguia, J., Huck, S., Oechssler, J., 2007. Imitation—theory and experimental evidence. *J. Econ. Theory* 136 (1), 217–235.
- Asch, S.E., 1952. Group forces in the modification and distortion of judgments. In: *Social Psychology*. Prentice-Hall, New York, pp. 450–501.
- Bao, T., Duffy, J., Hommes, C., 2013. Learning, forecasting and optimizing: an experimental study. *Eur. Econ. Rev.* 61, 186–204.
- Boylan, R.T., El-Gamal, M.A., 1993. Fictitious play: a statistical study of multiple economic experiments. *Games Econ. Behav.* 5 (2), 205–222.
- Brocas, I., Carrillo, J.D., Tarrasó, J., 2018. How long is a minute? *Games Econ. Behav.* 111, 305–322.
- Brown, G.W., 1951. Iterative solution of games by fictitious play. In: Koopmans, T.C. (Ed.), *Activity Analysis of Production and Allocation*. Chapman & Hall, London, pp. 374–376.
- Camerer, C., Hua Ho, T., 1999. Experience-weighted attraction learning in normal form games. *Econometrica* 67 (4), 827–874.
- Charness, G., Levin, D., 2005. When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *Am. Econ. Rev.* 95 (4), 1300–1309.
- Cournot, A.A., 1838. *Recherches sur les principes mathématiques de la théorie des richesses*. L. Hachette, Paris.
- Erev, I., Roth, A.E., 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *Am. Econ. Rev.*, 848–881.
- Frederick, S., Loewenstein, G., O'Donoghue, T., 2002. Time discounting and time preference: a critical review. *J. Econ. Lit.* 40 (2), 351–401.
- Friedman, D., Huck, S., Oprea, R., Weidenholzer, S., 2015. From imitation to collusion: long-run learning in a low-information environment. *J. Econ. Theory* 155, 185–205.
- Friedman, D., Oprea, R., 2012. A Continuous Dilemma. *Am. Econ. Rev.* 102 (1), 337–363.
- Fudenberg, D., Kreps, D.M., 1993. Learning mixed equilibria. *Games Econ. Behav.* 5 (3), 320–367.
- Goeree, J.K., Yariv, L., 2011. An experimental study of collective deliberation. *Econometrica* 79 (3), 893–921.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1 (1), 114–125.
- Huck, S., Leutgeb, J., Oprea, R., 2017. Payoff information hampers the evolution of cooperation. *Nat. Commun.* 8 (1), 15147.
- Huck, S., Normann, H.-T., Oechssler, J., 1999. Learning in Cournot oligopoly—an experiment. *Econ. J.* 109 (454), 80–95.
- Huck, S., Normann, H.-T., Oechssler, J., 2002. Stability of the Cournot process—experimental evidence. *Int. J. Game Theory* 31 (1), 123–136.
- Krajbich, I., Bartling, B., Hare, T., Fehr, E., 2015. Rethinking fast and slow based on a critique of reaction-time reverse inference. *Nat. Commun.* 6 (1), 7455.
- Nowak, M., Sigmund, K., 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's dilemma game. *Nature* 364 (6432), 56.
- Offerman, T., Potters, J., Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment. *Rev. Econ. Stud.* 69 (4), 973–997.
- Pettit, J., Friedman, D., Kephart, C., Oprea, R., 2014. Software for continuous game experiments. *Exp. Econ.* 17 (4), 631–648.
- Rand, D.G., Greene, J.D., Nowak, M.A., 2012. Spontaneous giving and calculated greed. *Nature* 489 (7416), 427.
- Rick, S., Weber, R.A., 2010. Meaningful learning and transfer of learning in games played repeatedly without feedback. *Games Econ. Behav.* 68 (2), 716–730.
- Rubinstein, A., 2007. Instinctive and cognitive reasoning: a study of response times. *Econ. J.* 117 (523), 1243–1259.
- Simon, H.A., 1972. Theories of bounded rationality. In: Radner, R., McGuire, C.B. (Eds.), *Decision and Organization*. North Holland, Amsterdam, pp. 161–176.
- Spiliopoulos, L., Ortmann, A., 2018. The bcd of response time analysis in experimental economics. *Exp. Econ.* 21 (2), 383–433.
- Vega-Redondo, F., 1997. The evolution of Walrasian behavior. *Econometrica* 65 (2), 375–384.