On the strategic value of ‘shooting yourself in the foot’: an experimental study of burning money

Michał Krol¹ · Magdalena Ewa Krol²

Accepted: 14 April 2019 / Published online: 26 April 2019
© The Author(s) 2019

Abstract

This paper uses a novel experimental design to investigate the strategic value of making a public payoff sacrifice (“burning money”) in a setting modelled after the battle-of-the-sexes game. Unlike prior studies, we find that subjects choose to burn money in a significant portion of decision trials and that burning makes the first movers more likely to achieve their preferred stage two equilibrium outcome. No such effect is observed in a control treatment where money burning is triggered unintentionally. This suggests that payoff sacrifices are important as deliberate signals of one’s further intentions, and can be a source of strategic advantage different from that stemming from being the first party to move. These results are supported by an analysis of subjects’ eye movements, which reveals patterns consistent with forward induction reasoning.

Keywords Forward induction · Burning money · First-mover advantage

1 Introduction

The classic game-theoretic concept of forward induction (Kohlberg and Mertens 1986) suggests that past actions of a player could signal to her counterparts what that player will do next. A number of experimental studies provided evidence consistent with this idea. In particular, people who sacrifice an ‘outside option’ or incur a cost in order to play a game with multiple equilibria have an increased chance of subsequently securing their preferred equilibrium outcome, reducing the risk of coordination failure

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s00182-019-00673-5>) contains supplementary material, which is available to authorized users.

✉ Michał Krol
michal.krol@manchester.ac.uk

Magdalena Ewa Krol
mkrol1@swps.edu.pl

¹ Department of Economics, The University of Manchester, Manchester, UK

² University of Social Sciences and Humanities, Wrocław, Poland

(Van Huyck et al. 1993; Cachon and Camerer 1996; Cooper et al. 1992, 1993; Brandts and Holt 1995).

More recently, Evdokimov and Rustichini (2016) collected answers to questions designed to elicit the players' reasoning in a battle-of-the-sexes game with an outside option. They found that what might limit the impact of forward induction reasoning is that players with an opportunity to move first are often unable to turn this into a source of strategic advantage, because even though they understand the signalling potential of their initial choice, they lack the confidence that their counterparts will understand the signal.

This result also sheds light on why, despite substantial evidence of the impact of outside options, little empirical support has so far been found for a related concept of 'burning money', whereby an option to sacrifice some payoffs before a game is played is said to be a source of advantage *whether the option is exercised or not* (Van Damme 1989; Ben-Porath and Dekel 1992; Hurkens 1996). In contrast with the theoretical argument, Beck et al. (2013) reported that having an option to burn money, exercised or not, does not increase one's payoffs in a 'credence game'. The authors found the lack of impact of a decision not to burn unsurprising because the underlying theoretical argument 'relies on iterated forward induction and thereby requires many layers of mutual knowledge of rationality'. However, they noted that a similar ineffectiveness of a decision to burn (which demands less iterative reasoning skills from the players in order to work) is puzzling. Blume et al. (2017) considered a scenario in which the sender of the signal can communicate to the counterpart not only its decision to burn money but also an explicit suggestion as to what it will do next. This made it easier for the receiver to understand the signal, and burning money had the desired effect. However, the frequency of burning was 'dramatically reduced' by even a slight increase of the associated cost (for other recent work on the impact of signalling costs on coordination in games, see Bilancini and Boncinelli (2018) or Masiliūnas (2017)).

These results indicate that the reason why people seem to use forward induction reasoning in outside option games but not in burning money games is that using it in the latter context is more complicated. Even though burning money (rather than just considering to do so) seems to send out a stronger and easier to understand signal of one's intentions, it carries an extra cost that might be too high given a limited degree of 'confidence in others'.

A way out of the problem is suggested by the recent work by Alekseev et al. (2017). They argue that the use of meaningful context and player labels in tasks requiring sophisticated strategic reasoning, e.g. games that rely on forward induction, is highly beneficial, as it can lead to choices being more consistent, knowledgeable and strategic. This is in line with earlier research (Chou et al. 2009) suggesting that presenting games in an abstract form of a payoff matrix, rather than in a simple, familiar context, can lead to subjects failing to behave strategically and reverting to basic heuristics. Indeed, in an early experimental study of burning money, Huck and Müller (2005) found that the outcome of the game is highly-dependent on how it is presented, with no strategic advantage gained by the first-mover under a normal form representation, and an outcome indicating a first-mover advantage but not forward induction when the game is presented in extensive form.

Motivated by these results, we use a novel experimental design to study forward induction. First, while we use a game similar to ‘battle-of-the-sexes’, we make it easier for the subjects to understand, describing it as a ‘conflict for resources’ visualized in a simplified graphic form instead of a payoff matrix, as in earlier studies. As an additional aid, we provide training sessions to teach subjects how to play the basic underlying game before introducing the option to burn money beforehand.

Similarly to Huck and Müller, we used a control treatment to disentangle forward induction from a more general first mover advantage stemming from being able to move first rather than simultaneously with one’s counterpart. Our control treatment differs slightly from that used by these authors, in that the first mover does not know which of her two available stage one actions will result in burning. Thus, the stage one choice can no longer signal stage two intentions through forward induction, while the timing of moves is preserved. Any gains of the first-mover relative to the control treatment can therefore be attributed to the fact that forward induction is possible in the main treatment, rather than to a change in timing and a generic first-mover advantage.

Finally, to gain more insight into the underlying reasoning processes, we investigate if subjects playing as receivers in the main experimental treatment (compared to those in the control treatment) will be more interested in acquiring information on whether the money was burned or not, and how acquiring it would determine their stage two actions. We do this by means of the eye-tracking technique, i.e. use a camera integrated with a computer system to identify the point on the screen that a subject is looking at any time. Existing research in cognitive science shows that, based on the way people visually examine information during problem-solving, it is possible to determine their cognitive traits like intelligence and working memory capacity (Hayes and Petrov 2016) or to reveal the objectives that they pursue (Borji and Itti 2014). In contrast with other cognitive process tracing techniques, like functional magnetic resonance imaging, eye-tracking makes it possible to study behavior under natural conditions, akin to a day-to-day computer task. It has been widely used not only to study bias in individual decisions (e.g. Król and Król 2019) but also specifically in experimental game theory (Polonio et al. 2015; Król and Król 2017; Stewart et al. 2016).

Unlike the previous literature, we find that subjects choose to burn money in a significant proportion (over a third) of all decision trials. Crucially, the likelihood of the first-mover/sender attaining its most preferred stage two equilibrium outcome is greater in the main treatment than in the control treatment, and greater when money was burned than when it was not. Furthermore, the effect of burning is greater in the main treatment than in the control treatment. Lastly, eye-tracking reveals that receivers are more likely to look at information on whether money was burned or not when burning is triggered intentionally by the sender. Furthermore, when the receiver is known to look at this information, choosing to burn has an impact on receivers’ behaviour that is more strongly beneficial for the sender.

Our contribution is therefore to provide a counterexample to existing studies, showing that *under certain conditions*, it is possible for burning money to become a source of strategic advantage, as postulated by game theory, to the extent that players would actually choose to bear the associated cost. In particular, in line with other recent studies, we point to the game being presented in simple form and context as a key factor allowing players to understand the game and behave strategically. We also demon-

strate the crucial role of the sender's deliberate intention to burn being clear to the counterpart, differentiating the impact of burning money from a first-mover advantage. Finally, we provide eye-tracking data consistent with forward induction reasoning.

2 Method

2.1 Subjects

The experiment was conducted at the University of Social Sciences and Humanities in Wrocław. A total of 96 subjects were recruited from the local population of undergraduate and postgraduate students; their average age was 22.5, and 50 of them were female. The protocol of the study was approved by the local Research Ethics Committee and the study was conducted in accordance with the Helsinki Declaration.

2.2 Procedure

The experiment was computerized—our stimulus presentation software was programmed in Embarcadero Delphi XE5, and the interface for remote communication with the eye-tracking workstation was programmed in C# using Microsoft Visual Studio Express. The experiment featured two experimental treatments (described in detail below), each comprising eight sessions with six subjects taking part in every session. Each subject was seated at a separate computer terminal and all were asked not to communicate with each other. Instructions were read aloud, shown on a board visible to all subjects, and displayed on each computer screen. One of the six computer terminals had attached underneath the screen a RED250 eye-tracking device, manufactured by SensoMotoric Instruments and set to 60 Hz frequency (thus, we recorded the eye-data of a total of $2 \times 8 \times 1 = 16$ subjects). The terminal had a 22 inch screen with resolution set to 1280×720 , and the distance between the subject's eyes and the screen was approx. 70 cm. For the subject seated at the eye-tracking terminal, we conducted a standard five-point semi-automatic calibration and validation procedure. Subjects are asked to look at small pulsating dots that appear on the screen in quick succession. After that, they are asked to look at a different set of stimuli to record the offset between the target and the gaze point identified by the eye-tracker (the average deviation was below 0.5° for all tested subjects). To detect eye-fixations, we used the SMI Vision Event Detector software with default settings (min. duration 80 ms, max. dispersion 100 px). The experiment took around 40 min and the average total payoffs were equivalent to approximately 10 USD in local currency (subjects received a fixed amount equivalent to 0.06 USD for every point scored in the game, with no possibility of negative winnings—see below).

2.3 Stimuli and design

Each session of the experiment consisted of three parts, the first two of which were intended to gradually familiarize subjects with the rules of the game played in part three, and were the same across the two treatments. Subjects played in pairs against

each other during a total of 32 rounds (split across the three parts of the study). A random matching procedure was designed to ensure that no subject played the same other subject more than once in the same player role and in the same variant of the game (see below).

2.4 Part I

The first part of each session consisted of five rounds. In each of them, subjects were shown two sets of geometric figures: a large set of seven figures (four circles and three squares), and a small set of five figures (four circles and one square). The position of the two sets on the screen was randomly chosen in each round: either the large set on the left and the small one on the right, or vice versa. Similarly, in every round subjects were randomly allocated the role of a 'circle player', represented by a round green smiley face icon, or a 'square player' (red square smiley face). The icons were initially shown in the centre of the screen, between the two sets, and the icon of the player whose role was assigned to the subject was highlighted in yellow (see Fig. 1).

In every round, the role of the other (not highlighted) player was randomly allocated to one of the other five subjects taking part in the session, but no two subjects would face each other more than once in the same role in this part of the experiment.

Each subject's task was to independently select either the large or the small set of figures, by dragging their player's icon to the left or to the right of the screen. Depending on their choices, each subject was then allocated a number of figures, and every figure converted to a fixed monetary amount. The figures were allocated according to the following rules:

Rule 1: If each player selected a different set, each gets all figures in her chosen set.

Rule 2: If both players selected the same set, the square player gets all squares in the chosen set, and the circle player gets all circles in the chosen set.

In order to help our subjects understand and implement these rules, they were given the option to 'simulate' the outcome, by dragging the icons of both players into hypothetical positions and observing the allocation of figures that would have followed prior to making a final decision (see an example in Fig. 2).

Once the subject finalized her choice by clicking a 'select and continue' button, the icon of the other player would start pulsating until the latter makes her own choice, and then move into the appropriate position, with the allocation of figures and the

Fig. 1 An example decision screen shown to subjects in part I of the study. Here, the large set is displayed on the left, and the subject was assigned the role of the square player (colour figure online)

Fig. 2 An example allocation of figures as shown to participants. Here, both player icons are on the left side of the screen, i.e. both have selected the set on the left (in this case this is the large set). Accordingly, the circle player receives four figures (marked in green), and the square player receives three figures (in red)

		square	
		large set	small set
circle	large set	4, 3	7, 5
	small set	5, 7	4, 1

Fig. 3 The normal form of the game played by subjects in part I of the experiment

subject's monetary payoffs updated accordingly. The subject could then reflect on the outcome and move on to the next round upon clicking a 'continue' button.

Note that the game in question is equivalent to the normal-form game presented in Fig. 3, and is similar to the classic 'battle-of-the-sexes' (BoS) game used by existing studies. In particular, choosing the large set is similar to selecting one's 'preferred outcome' in BoS. There are two pure-strategy Nash Equilibria, each with a different player choosing the large set ('preferred outcome'), and the other player choosing the small one (similar to the 'yield' strategy in BoS). There is also a mixed-strategy equilibrium in which each player chooses the preferred outcome/large set with a probability of $3/4$. While each of the two pure-strategy equilibria favors a different player, when choosing the same strategy both are worse off than in their less preferred equilibrium.

The difference between the two games is that the one used here is not completely symmetric. When both players choose the large set ('preferred outcome'), or when both choose the small one ('yield'), the circle player is better off than the square player. Importantly, the circle player receives four figures in either of these cases, the consequences of which will be discussed later.

2.5 Part II

The second part of every session consisted of 11 rounds. The only difference compared to part I was introducing a single new element to the game, black circles, with the following effect on payoffs:

Fig. 4 An example decision screen shown to participants in part II of the experiment, with two black circles in each set of figures

Rule 3: Any black circles allocated to the circle player do not count for the purpose of calculating her payoffs/monetary rewards.

There was always the same number of black circles in the large set and in the small set, which meant that the payoffs of the circle player were simply reduced by this number, while leaving the square player's payoffs unaffected. The number of black circles in each set was either one, two, or three, determined at random in each trial.

An example of this is shown in Fig. 4. In this example, choosing the large set (left) by both players would result in the circle player only getting the two circles there that are not coloured in black. Once these choices are finalized, the screen shown to the subject would resemble the one in Fig. 2, except for the presence of the black circles. In particular, the two black circles in the left set would not change colour to green, i.e. the circle player would only receive two points/figures instead of four. However, if the square player chose the small set (right), then all the five figures in that set would turn red and be allocated to her, while once again only the two white circles in the large set (plus the three squares) would turn green and be allocated to the circle player. All in all, the payoff matrix in this game is as shown in Fig. 3, except the payoffs of the circle player (first number in each cell) are reduced by the number of black circles (in this case, two).

2.6 Part III

The last, main part of each experimental session consisted of 16 rounds. In every round, the game was now played in two stages. First, the subject assigned the role of the circle player was shown two subgames, each similar to the games played in part II, and identical to each other except for the number of black circles in each set of figures which was larger (by one) in one of the two subgames than in the other. One of the two subgames was displayed in the upper half of the screen and the other in its bottom half, an example of which is shown in Fig. 5.

In every round, each pair of subjects was randomly assigned to one of two variants of the game: (V1) where the number of black circles in each set was one in one subgame and two in the other subgame; or (V2) where the number of black circles in each set was two in one subgame and three in the other subgame (thus, the three possible subgames across the two variants of the game corresponded to the three one-stage variants of the game used in part II). In addition, in every round, it was randomly determined

Fig. 5 An example decision screen shown to subjects in part III of the study. There is a single black circle in each set of figures in the ‘top’ sub-game, and two in each set in the ‘bottom’ one. The large set(s) of figures are shown on the right. The subject, assigned the role of the circle player, selected the ‘top’ sub-game and moved the player icons into hypothetical positions in both (note the black circle in the right set of the top subgame is hypothetically assigned to the square player, and so coloured in red) (colour figure online)

whether the subgame with the larger number of black circles would appear in the top or in the bottom half of the screen, and whether the large set of figures would appear on the left or right (in both subgames). In every round, each subject was also randomly assigned the role of circle or square, and randomly matched with one of the remaining five session participants. Similarly to the previous two parts of the experiment, the randomization procedure was designed to ensure that, throughout part III, no two players would play the same variant of the game against each other more than once in the same roles (subjects were made aware of this). Additionally, each subject would play the same number of times in each of the two roles (of circle and square).

In the first stage of the game, the task of the subject assigned the role of the circle player was to select the subgame to be subsequently played (at stage two) by the two players in the same way as they had done in the first two parts of the experiment. However, the way in which the circle player was to select the subgame varied with the experimental treatment in the way that we now describe.

2.7 Main treatment

In this case, the subject assigned the role of the circle player could directly choose one of the two subgames. Thus, by choosing the subgame with the lower number of black circles, the subject chooses not to burn any money, and by choosing the one with one additional black circle she chooses to burn one unit of payoffs (that is, in both variants of the game, V1 and V2, the player can either burn nothing at all or burn a single payoff point). The choice was submitted by either pressing the ‘A’ keyboard key to select the subgame shown at the top of the screen, or pressing ‘L’ to select the

one at the bottom. The selected subgame was shown as more opaque/highlighted (see Fig. 5). The subject could also drag-and-drop the player icons in both subgames in the usual manner, to see what would happen in each subgame if hypothetical set choices were made, prior to finalizing her choice of subgame. In the meantime, the square player would wait, being shown a blank screen with a pulsating message ‘waiting for the other player’.

Once the subgame was chosen by the circle player, i.e. at stage two, both subjects would be shown the two subgames, with the selected one now permanently highlighted (more opaque). The subjects could still drag-and-drop the player icons in both subgames, to see what would happen in each if hypothetical set choices were made. However, they were told that only their choice in the selected subgame would matter, as indicated by ‘choose left’ or ‘choose right’ caption on the button they had to click to finalize their set choices. The round would then end and payoffs would be allocated in the usual manner.

2.8 Solution by forward induction

Let us demonstrate how forward induction can be used to solve the game. Suppose first that, in the first variant of the game (V1, depicted in Fig. 5), the circle player chooses the subgame with one black circle. She is then guaranteed at least three figures, but can get more if the two players go on to choose different sets (which will happen with positive probability in any Nash Equilibrium). If she then selects the subgame with two black circles instead (i.e. chooses to ‘burn’ one figure/payoff point), this means that, by doing so, she is looking to improve on a guaranteed payoff of three figures. Hence, she must expect the pure strategy Nash Equilibrium to be played in which she selects the large set and the counterpart chooses the small one, as the other two equilibria cannot beat a payoff of three (in particular, the mixed-strategy Nash Equilibrium yields an expected payoff of $11/4$).¹ If the square player understands this, and expects circle to choose the large set, she will indeed want to select the small one. Thus, by choosing to ‘burn’ at stage one, the circle player can ensure a payoff of five figures. We will refer to this as ‘level one forward induction reasoning in the burning money game’, or FIR.1 for short. An analogous reasoning applies to variant (V2) of the game shown to subjects (choosing between two and three black circles).

Consider now the next step, which we will refer to as ‘level two forward induction reasoning in the burning money game’, or FIR.2 for short. Specifically, suppose that based on FIR.1 the players expect circle to end up with five figures when choosing to ‘burn’ at stage one, but instead circle opts not to burn. This would suggest that she expects to get more than five figures when choosing not to burn, and the only Nash

¹ Note how this means that we could not have used a symmetric game, which would require that the number of circles in the large set (call it x) equals the number of squares in that set, and the same is true for the number of circles/squares in the small set (y). It is then impossible to find a number of black circles (z) such that: (1) the lowest equilibrium payoff with burning ($2y - z$) does not exceed the guaranteed minimum payoff without burning (y); and (2) it is impossible for a subject to walk away with nothing ($y - z > 0$). Even if requirement (2) was relaxed to $y - z \geq 0$, we would still need $z = y$, which, among other things, would make it possible to infer if money was burned without looking at the other subgame (thus confounding the eye-tracking analysis—see the hypotheses section).

Equilibrium of the corresponding subgame in which this happens is again the one in which she selects the large set and the counterpart chooses the small one.

Thus, the circle player ends up with six figures by merely considering the option to burn money (once again, the same reasoning also applies to the second variant of the game). Note that we only make the distinction between FIR.1 and FIR.2 in the context of the specific game studied here, and do not claim that these are two generic types of forward induction reasoning that can be applied to other games.

What is important for our further considerations is that FIR.2, taking FIR.1 reasoning one step further, is likely to be more difficult for our subjects to conduct than FIR.1.

2.9 Control treatment

The control treatment differs from the main treatment in one feature only. Specifically, subjects assigned the role of the circle player in any given round of part III did not know, at stage one, which key ('A' or 'L') would select the subgame shown at the top of the screen, and which one would select the one at the bottom. They were told that this was decided beforehand by the experimenters by tossing a fair coin, once for every round of part III. If the result was 'heads', then pressing 'A' would select the game at the top of the screen, and pressing 'L' would select the one at the bottom. If the result was 'tails', the key roles would be reversed. In addition, subjects were told that the results of all the coin tosses had been written down on a sheet of paper placed next to them, which they would be able to inspect upon completion of the experiment. The change means that forward induction reasoning no longer applies, as the circle cannot signal her intentions via her stage one choice. At the same time, the relative timing of moves of the two players is preserved.

Thus, we believe that any first-mover advantage, defined as a payoff improvement resulting from moving first rather than simultaneously with the counterpart (independent of the nature of the initial move), would occur in the main treatment to the same extent as in the control treatment. Similarly, the Huck and Müller (2005) study used a control treatment where the first-mover 'selects' between two identical subgames. Although technically the player does not therefore make a payoff-relevant choice, i.e. is not a 'mover' in a strict sense, the authors still argue that 'the fact one player can make a first choice renders this player's preferences "focal" (even if the choice is materially irrelevant).' They refer to this as the 'physical timing hypothesis' and report that a first-mover advantage is present as a result. They also show that such 'materially irrelevant' physical timing is sufficient to induce the first-mover's preferred equilibrium in a similar way to what is observed in the burning money game (in which payoffs in one of the two subgames are different). In comparison, our treatments are designed so that each of them involves the same pair of subgames. This allows us to rule out the possibility that any observed differences between stage two outcomes across the treatments are due to payoffs being different in each case.

2.10 Hypotheses

Compared with existing studies, our experimental design includes a number of features that should make it easier for subjects to understand the game and carry out forward induction. First, presenting the game in a graphic form only requires subjects to internalize a small set of intuitive rules (described earlier) instead of analysing two payoff matrices of 8 numbers each, or one of 16 numbers, as in existing studies (a long-standing consensus in decision science is that graphical presentation of numerical information leads to better decision quality than tabular presentation, particularly when task complexity and information load are high (Remus 1987; Gettinger, Koeszegi, and Schoop 2012; Speier and Morris 2003). In the more specific context of experimental game theory, existing studies showed that presenting a game in an abstract form could lead to loss of experimental control and prevent players from acting strategically, while framing it in a familiar context and simple form can resolve the problem (Chou et al. 2009; Alekseev et al. 2017). In the same spirit, the two training parts of each experimental session should make it easier to understand the potential of burning money when it is introduced in part III.

Furthermore, allocating an easily-noticeable guaranteed payoff of between two and four (depending on the game variant) to the circle player (recall Fig. 1) facilitates putting a value on the option relinquished by the player through burning. This should simplify the forward induction process by making the game more similar to ‘outside option’ games in which forward induction reasoning is readily demonstrated. A potential downside is that it also creates a payoff asymmetry between the players, making it hard to judge the extent to which any observed changes in behaviour compared with existing studies are due to the simplified format vs. the asymmetric payoffs. Here, we try to control the latter effect as much as possible by interchangeably using variants of the game with different numbers of black circles, i.e. varying the average payoff of the circle player relative to that of the square player. The variant of the game is then included as a control variable in the estimated regressions to verify that any effects of burning money that we observe are not mere artefacts of the asymmetry between the players.

Overall, with the changes described above, we expect the generally complex forward induction reasoning to become more accessible to subjects relative to simple heuristics based on physical timing (such as ‘the player who moves first should get the large set’). Thus, we should be able to demonstrate any benefits from having an opportunity to burn money that are additional to and separate from the physical timing effects. Specifically, we hypothesize that, compared with the control treatment, the outcome of part III of the main treatment will be closer to the circle player’s preferred stage two equilibrium, with the circle player more likely to choose the large set, and the square player more likely to choose the small one. Since FIR.1 is easier to carry out than FIR.2, we also expect the difference between the treatments to be more pronounced when money is burned at stage one than when it is not.

With regards to the eye-tracking data, we build on the fact that forward induction reasoning is based on drawing conclusions from what exactly the first-mover does at stage one, whereas all that matters for a first-mover advantage is the very fact of

Fig. 6 The main AOI, enclosed by red dotted lines. This is the part of the screen the subject assigned the role of the square player would need to look at to compare the number of black circles between the two subgames. The AOI covered approximately 10% of the screen (colour figure online)

moving first. Thus, making the effort to find out what the counterpart did at stage one should be a signature of forward induction reasoning. To measure the intensity of that effort, we record all eye fixations located in the area depicted in Fig. 6 (a fixation is an act of pausing one's gaze on any part of the visual field).

This 'Area-of-Interest' (AOI, in the eye-tracking terminology) is specified as the set of two equally sized rectangles, enclosing the two collections of circles in part III, stage two choice screen of the square player in the subgame that has NOT been selected by the counterpart. This is where the subject allocated the role of square must look in order to find out if the circle player burned money at stage one (recall that we varied the number of black circles in each subgame between the trials, so the subject could not otherwise be certain if the selected subgame has more or less black circles than the other subgame).

We hypothesized that subjects would be more likely to look at the AOI in the main treatment than in the control treatment (as a signature of more forward induction reasoning taking place). Furthermore, in trials in which the square players look at the AOI the differences between the treatments postulated above should be particularly strong.

3 Results and discussion

3.1 Comparisons of part III of the main and control treatments with part II

We begin by reporting the frequencies of the $2 \times 2 =$ four possible stage two outcomes of the game, depending on the experimental treatment (main vs. control), and the subgame selection triggered (intentionally or not, depending on treatment) by the circle player: not burn vs. burn. For reference, we include the frequencies obtained in part II (which is when the black circles are introduced but players still move simultaneously). Finally, we separately report the outcome frequencies for those trials in which the subject playing as the square player was seated at the eye-tracking terminal and did or did not look at the AOI representing the subgame that has not been selected by the counterpart. The results are presented in Table 1.

The first thing we may wish to note from Table 1 is that the frequency of the circle player’s preferred outcome (where the circle player selects the big set and the square player chooses the small one) was greater in the main treatment than in part II of the game, but no such change occurred in the control treatment. Figure 7, showing the evolution of the said frequency over the course of part II and part III, suggests no obvious time trend within a given part of the study/experimental treatment, while illustrating the above differences between them.

To verify the statistical significance of the differences observed above, we estimated a mixed-effects binary probit model with random intercept and slope effects grouped by experimental session (to allow for within-experimental session correlated errors). The binary dependent variable took a value 1 if the circle player’s preferred outcome occurred in a given trial (each part II and part III trial constituted a single observation), and zero otherwise. We included two binary independent variables—‘control’ and ‘main’, each taking a value 1 if the trial occurred during part III of the control and main treatment respectively (the other possibility, set up as the regression’s reference category, was that the trial occurred during part II). Additionally, we included a control variable representing the number of black circles in each set (for part III trials, this was based on the subgame that was selected at stage one, and, for all trials, we subtracted two from each value, i.e. used the variant of the game with two black circles per set as the reference value). The resulting fixed effect coefficient estimates are reported in Table 2.

The results indicate that the number of black circles per set has no effect on the likelihood of the circle player obtaining her preferred outcome in part II ($\beta_{n\text{-black}} = -0.144$, $p = 0.170$), and there is no significant difference in this respect between part II and part III of the control treatment ($\beta_{\text{control}} = -0.128$, $p = 0.333$, and $\beta_{n\text{-black} \times \text{control}} = -0.296$, $p = 0.101$). However, in part III of the main treatment, the said likelihood is higher than in part II given a number of black circles per set equal to the reference value of two ($\beta_{\text{main}} = 0.608$, $p < 0.001$), and the number of black circles has a significantly more positive effect on the likelihood of the circle player’s preferred outcome than in part II ($\beta_{n\text{-black} \times \text{main}} = 0.770$, $p < 0.001$). We may summarize this as follows.

Table 1 The frequencies (%) of the four possible stage two outcomes depending on the experimental treatment, the initial decision of the circle player, and whether or not the square player seated at the eye-tracking terminal looked at the subgame not chosen by the circle player (note that the ‘overall’ frequencies also include cases where the square player was not seated at the eye-tracking terminal)

			Circle small Square Small	Circle small Square big	Circle big Square Small	Circle big Square big	N
Part II			8.14	41.48	8.14	42.23	528
Main treatment (intentional burning)	Overall		10.16	36.46	21.61	31.77	384
	Circle NOT burn	Overall	12.08	50.	11.67	26.25	240
		Square not look	25	37.5	25	12.5	8
		Square look	7.41	48.15	3.7	40.74	27
	Circle burn	Overall	6.94	13.89	38.19	40.97	144
		Square not look	0	31.25	18.75	50	16
		Square look	7.69	7.69	69.23	15.38	13
	Overall		6.51	58.85	7.03	27.6	384
	Circle NOT burn	Overall	4.74	60.53	7.89	26.84	190
		Square not look	3.45	62.07	10.34	24.14	29
		Square look	0	33.33	33.33	33.33	3
Control treatment (unintentional burning)	Circle burn	Overall	8.25	57.22	6.19	28.35	194
		Square not look	3.85	53.85	15.38	26.92	26
		Square look	0	66.67	16.67	16.67	6
	Overall		6.51	58.85	7.03	27.6	384

Result 1 The frequency of the circle player’s preferred outcome is significantly greater in part III of the main treatment than in part II, particularly when the number of black circles per set is large. However, the same is not true for the control treatment.

In contrast with the existing literature, this result suggests that forward induction can play a role that is distinct from the effect of physical timing/first mover advantage. In particular, the first mover’s preferred equilibrium is more likely to occur under dynamic (rather than simultaneous) play only in the treatment in which forward induction reasoning is possible, while no difference from simultaneous play is observed in the control treatment.

Fig. 7 The evolution of the frequency of occurrence of the circle player's preferred outcome (circle big, square small), over the course of part II and part III (for each of the two experimental treatments)

Table 2 Coefficient estimates of a mixed-effects binary probit model of the probability of the circle player's preferred outcome occurring ('circle big, square small'), modelled as a function of the experimental treatment and the number of black circles

Probability of the circle player's preferred outcome (part II and part III data combined) N = 1296						
	β	SE	z	p	95% CI	
					Lower	Upper
Intercept	- 1.415	0.102	- 13.887	<0.001**	- 1.631	- 1.200
n-black	- 0.144	0.104	- 1.386	0.170	- 0.342	0.056
Control	- 0.128	0.132	- 0.967	0.333	- 0.386	0.132
n-black $\times$ control	- 0.296	0.180	- 1.646	0.101	- 0.643	0.053
Main	0.608	0.110	5.507	<0.001**	0.398	0.820
n-black $\times$ main	0.770	0.154	4.989	<0.001**	0.467	1.065

Note. Part II of the game is the reference category, i.e. both treatment dummy variables are equal to zero when the given instance of the game took place during part II (**significant at $p < 0.05$). Additionally, for part III trials, the number of black circles is based on the subgame that was selected at stage one, and we subtracted two from each value, i.e. centered the variable around a baseline value of 2

Additionally, results of the Pearson correlation also indicated that there was no significant association between the total payoff earned by a player during part I and part II and the player's behaviour during part III of the main treatment, specifically the frequency of the player burning money ($r = -0.10$, $p = 0.52$), of choosing the big set when playing as circle ($r = 0.06$, $p = 0.69$), or of choosing the big set when playing as square ($r = 0.18$, $p = 0.25$). Thus, there is no reason to suspect that the different patterns observed during part III were caused by different histories of play during the first two parts of the study.

We therefore proceed to focus on part III, specifically on the differences between the experimental treatments, as well as on the effect of the stage one behaviour of the

circle player on the players' subsequent decisions at stage two. In other words, having established that the likelihood of the first mover's preferred outcome depends on the treatment, we now investigate in more detail the reasons why this is the case.

3.2 The propensity to burn money

Result 2 The overall frequency of burning money in the main treatment was 37.5%. This is significantly different from 0 based on a one-sample one-tailed non-parametric Wilcoxon-test [$Z = 2.52$, $p = 0.007$].

The frequency of burning of almost 40% of the trials of the main treatment is both significantly above zero and higher than the very small fractions reported by previous studies (e.g. 6% in Huck and Müller 2005). One reason for this could be the reduction in the complexity of forward induction reasoning brought about by the novel features of our experimental design (graphical presentation, training sessions, and the easily-noticeable guaranteed payoff of the circle player). This made the circle players more likely to use money-burning, as more of them understood its potential and thought it more likely that the counterparts would understand the signal. In addition, we see in Table 1 that the outcome following no opportunity to burn (i.e. in part II) was worse for the circle player than what is typically obtained in 'battle-of-the-sexes' games (e.g. in Cooper et al. (1993) both players selected the preferred outcome with a probability of around 2/3). The one exception from this was when money was burned in the main treatment, in which case the outcome was almost the exact opposite of what was observed in the training sessions. In other words, the circle players started from a more disadvantaged position than in the canonical 'battle-of-the-sexes' games, thus making it more imperative for them to look for ways to improve it. As burning seemed the only way in which they could do so, they often opted for this strategy despite the associated self-inflicted cost. At the same time, in the model we estimate next, we will see that making the circle player generally worse or better off relative to the square player, i.e. manipulating the overall payoff asymmetry between them, had no significant effect on the players' behaviour.

3.3 The effect of burning money and evidence of FIR.1

To determine the statistical significance of the variance, observed in Table 1, in the frequency of choosing the big set by subjects when assigned the role of the circle player (particularly across treatments and depending on whether money was burned or not), once again we estimated a mixed-effects binary probit model. The model includes random intercept and slope effects nested by subject and experimental session, to allow for within subject/experimental session correlated errors. This allows us to control both the subject- and group-specific variation stemming from pairwise interactions between a specific set of subjects. Each trial belonging to part III of the study constitutes a single observation. The dependent variable is a binary variable taking a value 1 if the subject assigned the role of circle chose the big set, and 0 otherwise. The three independent binary variables are: 'burn' (taking a value 1 if money was burned), 'control' (taking

Table 3 Coefficient estimates of a mixed-effects binary probit model of the probability of choosing the big set by the circle player modelled as a function of the variant of the game, money-burning, and treatment (**significant at $p < 0.05$)Probability of choosing the big set by the circle player (based on part III data) $N = 768$

	β	SE	z	p	95% CI	
					Lower	Upper
Intercept	- 0.353	0.111	- 3.195	<0.001**	- 0.571	- 0.136
Game-variant	0.074	0.104	0.706	0.481	- 0.131	0.278
Burn	1.167	0.157	7.416	<0.001**	0.858	1.475
Control	- 0.103	0.128	- 0.805	0.421	- 0.354	0.148
Burn $\times$ control	- 1.143	0.203	- 5.623	<0.001**	- 1.542	- 0.744

Note. The reason why we used the main treatment (rather than the control treatment) as the reference category was that the ‘burn’ coefficient then estimates the effect of burning in the main treatment, which was our main focus

a value 1 if the trial belonged to the control treatment), and ‘game-variant’ (taking a value 1 if the two subgames contained respectively 2 and 3 black circles, and 0 if they contained 1 and 2 black circles). The results are presented in Table 3.

We find that, in the main treatment, having chosen to burn increases the likelihood that the first mover will subsequently select the big set ($\beta_{\text{burn}} = 1.167$, $p < 0.001$), but that this effect is significantly weakened in the control treatment ($\beta_{\text{burn} \times \text{control}} = -1.143$, $p < 0.001$).

We modelled the frequency of choosing the big set by subjects when assigned the role of the square player in a similar fashion (once again, nesting the random effects by subject and experimental session). The only difference was that, in this case, we included data from the eye-tracking-equipped terminal (one in each session) indicating whether or not the subject playing as square fixated on the AOI depicted in Fig. 6 during decision time. We included two binary independent variables: ‘look’ (taking a value 1 if the subject did look at the AOI) and ‘not-look’ (taking a value 1 if the subject did not look at the AOI). If the subject was seated at a terminal without eye-tracking, both variables would take a value of 0. In the current section, we focus on the behavioral results from terminals not equipped with eye-tracking (look = not-look = 0). This corresponds to the first five coefficient estimates in Table 4 (above the dashed line), and is based on the same set of predictors as the model reported in Table 3.

In a mirror image of the results in Table 3, we find that, in the main treatment, the first-mover’s decision to burn decreases the likelihood that the second mover (square) will subsequently select the big set ($\beta_{\text{burn}} = -0.596$, $p < 0.001$), but that this effect is significantly weakened in the control treatment ($\beta_{\text{burn} \times \text{control}} = 0.504$, $p = 0.035$).

Thus, the results in this section are consistent with FIR.1, and may be summarized as follows.

Result 3 Compared with choosing not to burn, burning money increases the likelihood of the first mover choosing the big set, and of the second mover choosing the small

Table 4 Coefficient estimates of a mixed-effects binary probit model of the probability of choosing the big set by the square player modelled as a function of the variant of the game, money-burning, treatment, and whether or not the subject looked at the specified AOI (**significant at $p < 0.05$; *significant at $p < 0.10$)

Probability of choosing the big set by the square player (based on part III data) $N = 768$

	β	SE	z	p	95% CI	
					Lower	Upper
Intercept	0.687	0.121	5.699	<0.001**	0.451	0.924
Game-variant	0.053	0.110	0.482	0.630	- 0.162	0.268
Burn	- 0.596	0.160	- 3.725	<0.001**	- 0.910	- 0.282
Control	0.488	0.163	2.970	0.003**	0.168	0.808
Burn $\times$ control	0.504	0.239	2.109	0.035**	0.035	0.972
Look	0.580	0.362	1.605	0.109	- 0.129	1.290
Not-look	- 0.659	0.476	- 1.385	0.167	- 1.594	0.275
Burn $\times$ look	- 1.483	0.552	- 2.684	0.007**	- 2.567	- 0.398
Burn $\times$ not-look	1.448	0.616	2.351	0.019**	0.239	2.657
Control $\times$ look	- 1.274	0.853	- 1.494	0.135	- 2.948	0.400
Control $\times$ not-look	0.544	0.568	0.958	0.338	- 0.570	1.658
Burn $\times$ control $\times$ look	2.085	1.149	1.816	0.070*	- 0.169	4.340
Burn $\times$ control $\times$ not-look	- 1.553	0.760	- 2.044	0.041**	- 3.045	- 0.062

Note. The binary variables 'look' and 'not-look' are both included as they are not mutually exhaustive (the third possibility is that the subject is not seated at the eye-tracking terminal and both variables are equal to zero)

set, but only in the main treatment in which forward induction reasoning is possible and burning can signal the first mover's subsequent intention.

3.4 The effect of the mere ability to burn money and evidence of FIR.2

When money is not burned (burn = 0), the square player is more likely to choose the big set in the control treatment than in the main treatment (Table 4, $\beta_{\text{control}} = 0.488$, $p = 0.003$). However, the analogous (opposite) effect of the treatment on the circle player's behaviour is not statistically significant (Table 3, $\beta_{\text{control}} = - 0.103$, $p = 0.421$). Looking at the data in Table 1, we see that, when money is not burned, the frequency of the first mover's most preferred outcome (circle big, square small) is higher in the main treatment (11.67%) than in the control treatment (7.89%), while the frequency of her least preferred outcome (circle small, square big) is lower in the main treatment (50%) than in the control treatment (60.53%; note that the 'overall' figures in Table 1 are based on all data, while the regression estimates we considered so far are based on excluding data from the single eye-tracking terminal).

Thus, there is some evidence to suggest that, consistent with FIR.2, even choosing not to burn could be a source of a strategic advantage, because the outcome is then more favourable to the first mover than when the same subgame is selected at random and cannot signal the subsequent intention. In other words, the mere opportunity to

intentionally burn money, even if not acted upon, could give a player a gain additional to any first-mover advantage that stems solely from the fact of being the first player to move.

At the same time, we should note that the frequency of the first mover's least preferred outcome in part II (Table 1, 41.48%) is even smaller than when money is not burned in the main treatment (50%). This could be caused by the relative timing of part II and part III trials, but could also reflect a paradox inherent in FIR.2. Specifically, if, as evidence suggests, subjects show a tendency to treat intentional burning as a signal of the first-mover's intention to choose the big set, then they might analogously treat not choosing to burn as a signal of an opposite intention, thus leading to a worse outcome for the first-mover than when no signal is received (as in part II).

Overall, our results support FIR.1 more strongly than FIR.2. As we have already noted, and as has also been observed by recent literature, FIR.1 requires less iterative reasoning capability on behalf of the players than FIR.2, i.e. it is easier to understand the strategic implications of the counterpart choosing to burn than of the counterpart choosing not to do so (which may be taken as simple avoidance of an unnecessary cost). Our results are consistent with this idea, since we observe that burning influences the counterpart more than choosing not to do so. This also further supports the notion that burning money can be a source of strategic advantage separate from and additional to any physical timing effects. In particular, if moving first was the only source of circle player's advantage, then there is clearly no reason for the square player to respond differently depending on circle's initial move, since either way the circle player has moved first.

3.5 Eye-tracking analysis

The behavioural data which we collected and analysed above has already provided support for our main hypothesis: that burning money can influence future actions in a way distinct from a first-mover advantage. However, it is still interesting to consider eye-tracking as an auxiliary tool that can further support, and perhaps refine, the insights from behavioural observations.

In particular, the frequency of having looked at the specified AOI (the subgame that was not selected) by subjects playing as square and seated at the eye-tracking terminal in the control treatment (look = control = 1) was 14.1%, and was significantly smaller than the same frequency in the main treatment, equal to 62.5% [Mann–Whitney $U = 63$, two-tailed $p = 0.001$].²

Result 4 The second-movers are more likely to attend to information about the subgame that was not selected at stage one when the subgame is chosen intentionally, rather than randomly, by the first-mover.

² This may be compared to the same frequency calculated for the circle players, which equalled 14.5% in the main treatment and 13.8% in the control treatment. Clearly, regardless of the treatment the circle player has no reason to look at the subgame that was not selected, having already familiarized herself with both subgames prior to triggering the subgame selection. Thus, it might potentially be argued that the square players in the control treatment are as uninterested in the AOI as the circle players are in both treatments.

Clearly, if moving first (whatever the exact decision) was the only source of the circle player's strategic advantage, then there is no reason for the square player to check if money was burned or not, and no reason to do so more often when the subgame selection is certain to be in line with the circle player's intentions. Thus, this result further supports the idea that prior choices can be a source of a strategic gain by signalling future intentions, in a way distinct from a first-mover advantage.

Additionally, it is important to verify that looking at the specified AOI is a sign of acquiring information about circle player's past move for the purpose of informing one's own strategy (rather than, for example, mere curiosity). In particular, to check if the fact of looking at the AOI was related to the observed choices of the square players, we consider the remaining, bottom part of Table 4 (below the dashed line). Compared with the reference category of subjects who were not seated at the eye-tracking terminal (look = not-look = 0), and who therefore might have looked at the AOI or not, subjects who are *known* to have looked at the AOI are more strongly influenced by burning in the main treatment ($\beta_{\text{burn} \times \text{look}} = -1.483$, $p = 0.007$). That is, the likelihood that they would choose the big set decreases more as a result of burning than for subjects in the reference category. Conversely, subjects who are known to have *avoided* the AOI are *less* strongly influenced by burning in the main treatment ($\beta_{\text{burn} \times \text{not-look}} = 1.448$, $p = 0.019$). This tendency can also be seen in Table 1: for instance, in the main treatment, the frequency of the circle player's most preferred outcome increases from 11.67% to 38.19% across all subjects when money is burned, but the change is considerably larger, from 3.7% to 69.23%, for subjects known to have looked at the AOI.

At the same time, the above effect of looking/not looking at the AOI in the main treatment is to an extent offset in the control treatment ($\beta_{\text{burn} \times \text{control} \times \text{not-look}} = -1.553$, $p = 0.041$; and $\beta_{\text{burn} \times \text{control} \times \text{look}} = 2.085$, $p = 0.070$). In other words, looking at the subgame that was not selected is less important for the second-mover's subsequent actions when the subgame selection is determined randomly rather than being made according to the first-mover's intention. In Table 1, this is to some degree reflected in the fact that, in the control treatment, the frequencies in the look/not look cells tend to trace the corresponding overall frequencies more closely than in the main treatment.³

Result 5 The effect of burning money (Result 3) is stronger in the main treatment (and stronger relative to the control treatment) for subjects who look at the subgame that was not selected than for those who do not.

This result further suggests that subjects' behaviour is driven by FIR.1 rather than FIR.2. If subjects were motivated by FIR.2, then they might have been inclined to consider the subgame that was not selected in order to carry out the full forward induction process, and to do so more often in the main treatment than in the control treatment (as per Result 4). However, in that case looking at the AOI would not have been associated with a greater change in behaviour depending on whether money was

³ However, we should note that, in case of the control treatment, those frequencies should be interpreted with caution, because subjects looked at the AOI in only 14.1% of the trials. Combined with the fact that only one subject in each session was eye-tracked, and that some of the four combinations of choices were rarely observed, this means that frequencies in some of the Table 1 cells represent a very small number of observations and are not very reliable. The advantage of using the mixed regression model instead of raw frequencies is that it takes this into account when estimating the significance of the key comparisons.

burned or not, as the mere consideration of burning is sufficient in FIR.2 and the exact action does not matter. Result 5 could therefore be seen as more evidence in favour of FIR.1 but against FIR.2.

Results 4 and 5 further underline that, to be successful, burning must actually occur, be intentional, and noticed by the other party. In particular, the counterpart must be willing to pay attention not only to the current state of play but also to ‘what might have been’.

3.6 Possible extensions

Since our results underline the role of the ease with which the counterpart can read the first-mover’s intention, one might wonder how the results would change if the circle player could convey an explicit ‘cheap talk’ message of her further intention together with the initial decision (as in Blume et al. 2017). If this could only be done given a decision to burn, then there is little doubt that an even greater proportion of the players would do so, and the effect on the stage two outcome would be even stronger. However, if the circle player could, alternatively, choose not to burn but indicate an intention to choose the big set, then this would possibly shift the stage two outcome somewhat towards the circle player’s preferred equilibrium, although not to the same extent as the same suggestion accompanied by burning (based on our results, we believe burning would make the message more credible). In other words, both options available to the circle player would become more effective, and the overall effect of cheap talk on the frequency of burning would depend on the relative magnitude of these changes.

One could also introduce a direct monetary cost of acquiring information about the first-mover’s initial choice (additional to the existing mental cost of allocating attentional resources to the subgame that was not chosen). Apart from being an interesting new scenario to study in future research, this could also be seen as a potential alternative to eye-tracking, bypassing the existing hardware resource constraints of only being able to monitor the gaze of one of the subjects in each session. In particular, by manipulating the cost of acquiring information, one could indirectly manipulate the likelihood that subjects would do so, while also limiting the noise due to accidental or curiosity-driven looking at the AOI. We anticipate that increasing the cost of acquiring information would strengthen the effect of burning on those subjects who choose to acquire it, but might also reduce the first-movers’ willingness to burn, as they would face an increased possibility that the counterpart will not notice this, opting to avoid the cost of acquiring the information.

4 Conclusions

In this paper, we reported the results of an experiment in which subjects play a two player game with multiple equilibria. Game theory has long suggested that a player could benefit from an option to carry out a public payoff sacrifice beforehand, which might convince the rival of the player’s intention to play consistently with the first-mover’s preferred equilibrium, thus inducing the rival to do likewise. However, so far

no experimental studies have proved that ‘shooting yourself in the foot’ could lead to strategic benefits that would be high enough for players to choose to bear the cost, and distinct from whatever is gained by simply being the first player to move.

We attributed this lack of existing evidence to the complexity of the forward induction reasoning required to understand the potential of burning money, and used a novel experimental design that simplifies the problem when seen from the subjects’ point of view. As a result, we found that money-burning does have the effect postulated by theory, making the first-mover more likely to attain the desired equilibrium outcome. These benefits from burning were greater than in a control treatment in which the second-mover did not know if the counterpart actually had intended for the money to burn, suggesting that burning conveys one’s subsequent intentions and provides one with a leverage additional to any first-mover advantage.

We also analysed the players’ eye-movements accompanying the observed decisions. This revealed that subjects make the effort to seek out information about their counterparts’ *intentional* past moves, and use it to guide their own subsequent choices.

Overall, the results emphasize that what matters for the outcome of strategic interactions is not only the current state of play (as per the well-known subgame consistency principle) but also how the parties arrived at the present situation and what intentions and knowledge they had along the way.

Funding Funding was provided by Polish Ministry of Science and Higher Education (BST/Wroc/2016/A/6).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alekseev A, Charness G, Gneezy U (2017) Experimental methods: when and why contextual instructions are important. *J Econ Behav Organ* 134:48–59
- Beck A, Kerschbamer R, Qiu J, Sutter M (2013) Shaping beliefs in experimental markets for expert services: guilt aversion and the impact of promises and money-burning options. *Games Econ Behav* 81(1):145–164
- Ben-Porath E, Dekel E (1992) Signaling future actions and the potential for sacrifice. *J Econ Theory* 57(1):36–51
- Bilancini E, Boncinelli L (2018) Signaling with costly acquisition of signals. *J Econ Behav Organ* 145:141–150
- Blume A, Kriss PH, Weber RA (2017) Pre-play communication with forgone costly messages: experimental evidence on forward induction. *Exp Econ* 20(2):368–395
- Borji A, Itti L (2014) Defending yarbus: eye movements reveal observers’ task. *J Vis* 14(3):29
- Brandts J, Holt CA (1995) Limitations of dominance and forward induction: experimental evidence. *Econ Lett* 49(4):391–395
- Cachon GP, Camerer CF (1996) Loss-avoidance and forward induction in experimental coordination games. *Q J Econ JSTOR* 111:165–194
- Chou E, McConnell M, Nagel R, Plott CR (2009) The control of game form recognition in experiments: understanding dominant strategy failures in a simple two person ‘Guessing’ game. *Exp Econ* 12(2):159–179
- Cooper R, DeJong DV, Forsythe R, Ross TW (1992) Forward induction in coordination games. *Econ Lett* 40(2):167–172

- Cooper R, DeJong DV, Forsythe R, Ross TW (1993) Forward induction in the battle-of-the-sexes games. *Am Econ Rev JSTOR* 83(5):1303–1316
- Evdokimov P, Rustichini A (2016) Forward induction: thinking and behavior. *J Econ Behav Organ* 128:195–208
- Gettinger J, Koeszegi ST, Schoop M (2012) Shall we dance? The effect of information presentations on negotiation processes and outcomes. *Decis Support Syst* 53(1):161–174
- Hayes TR, Petrov AA (2016) Mapping and correcting the influence of gaze position on pupil size measurements. *Behav Res Methods* 48(2):510–527
- Huck S, Müller W (2005) Burning money and (pseudo) first-mover advantages: an experimental study on forward induction. *Games Econ Behav* 51(1):109–127
- Hurkens S (1996) Multi-sided pre-play communication by burning money. *J Econ Theory* 69(1):186–197
- Kohlberg E, Mertens J-F (1986) On the strategic stability of equilibria. *Econometrica* 54(5):1003
- Król M, Król ME (2017) A novel approach to studying strategic decisions with eye-tracking and machine learning. *Judgm Decis Mak* 12(6):596–609
- Król M, Król ME (2019) Inferiority, not similarity of the decoy to target, is what drives the transfer of attention underlying the attraction effect: evidence from an eye-tracking study with real choices. *J Neurosci Psychol Econ*. <https://doi.org/10.1037/npe0000104>
- Masilūnas A (2017) Overcoming coordination failure in a critical mass game: strategic motives and action disclosure. *J Econ Behav Organ* 139:214–251
- Polonio L, Di Guida S, Coricelli G (2015) Strategic sophistication and attention in games: an eye-tracking study. *Games Econ Behav* 94:80–96
- Remus W (1987) A study of graphical and tabular displays and their interaction with environmental complexity. *Manag Sci* 33(9):1200–1204
- Speier C, Morris MG (2003) The influence of query interface design on decision-making performance. *MIS Q* 27(3):397–423
- Stewart N, Gächter S, Noguchi T, Mullett TL (2016) Eye movements in strategic choice. *J Behav Decis Mak* 29(2–3):137–156
- Van Damme E (1989) Stable equilibria and forward induction. *J Econ Theory* 48(2):476–496
- Van Huyck JB, Battalio RC, Beil RO (1993) Asset markets as an equilibrium selection mechanism: coordination failure, game form auctions, and tacit communication. *Games Econ Behav* 5(3):485–504

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.