

Sadanand, Asha B.

Article

Ideal reactive equilibrium

Games

Provided in Cooperation with:

MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Sadanand, Asha B. (2019) : Ideal reactive equilibrium, Games, ISSN 2073-4336, MDPI, Basel, Vol. 10, Iss. 2, pp. 1-18, <https://doi.org/10.3390/g10020019>

This Version is available at:

<https://hdl.handle.net/10419/219242>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>

Article

Ideal Reactive Equilibrium

Asha Sadanand

Department of Economics, University of Guelph, 50 Stone Road East, Guelph, ON N1G 2W1, Canada; asadanand@uoguelph.ca; Tel.: +1-(519)-824-4120 (ext. 58947)

Received: 28 December 2018; Accepted: 4 April 2019; Published: 15 April 2019

Abstract: Refinements of the Nash equilibrium have followed the strategy of extending the idea of subgame perfection to incomplete information games. This has been achieved by appropriately restricting beliefs at unreached information sets. Each new refinement gives stricter and more mathematically-complicated limitations on permitted beliefs. A simpler approach is taken here, where the whole idea of beliefs is dispensed with, and a new equilibrium concept, called the ideal reactive equilibrium, that builds on some pioneering work by Amershi, Sadanand and Sadanand on thought process dynamics, is developed.

Keywords: refinements of Nash equilibrium; dynamic games; imperfect information; beliefs

1. Introduction and Literature

In the last several decades, much attention has been devoted to refinements. Tremendous progress has been made in identifying ways to refine equilibria and establish which are the most sensible and should occupy our focus. Each round of refinements has improved upon previous work, typically requiring more difficult constructs involving formal mathematics. Explorations into beliefs, many of them mathematically mind bending, was aimed at determining which were reasonable. Reasonableness was defined in several ways, from limiting beliefs arising from a sequence of vanishing perturbation games, to restrictions arising from the construction of counterfactual beliefs when players are surprised by unexpected behaviour. While these developments have given us much insight into the workings of games with many subtle aspects to appreciate, there is also value in returning to some simplicity. This is the approach taken here. Our focus will be on dynamic games only, that is games played out in time, where the timing of moves is given by the extensive form.

All the complexity started with the desire to extend subgame perfection to dynamic games without full information. For an extensive-form game, Nash equilibrium and any of its standard refinements require that each player give an entire plan of action that specifies what a player would choose in every possible instance where the player could be asked to move. In addition, with incomplete information, many recent equilibrium concepts require construction of consistent beliefs that support the chosen actions on unreached information sets. In contrast, novice beginners to game theory often just specify the equilibrium path and do not see the importance of indicating the rest of the choices. Specifying these off-equilibrium path choices at information sets and justifying the choices by specifying beliefs were a way to operationalize the notion of rational choices at each information set. This was thought to be very useful because it would eliminate “incredible threats” in games without full information.

As I will argue here, this motivation to operationalize subgame perfection in dynamic games of incomplete information resulted in an over-exuberance of exotic belief construction, beyond what is needed to eliminate incredible threats and to achieve the dynamic consistency of strategies. Instead, I will return to the basics, with perhaps even a seemingly novice approach, to construct a new equilibrium concept. In some important work on dynamic games by Amershi, Sadanand and Sadanand [1–4], some basic ideas on thought process dynamics have been established. Here, I start with that thought process and substantially extend it in formulating a new equilibrium, which I call the

Ideal Reactive Equilibrium (IRE). This new equilibrium concept applies to dynamic games only, that is games played sequentially in time, where the extensive form gives the actual order of the moves.

In terms of literature,¹ very briefly, there are two genres of refinements: the trembles' refinements and the forward induction (deviations) refinements.² As the focus here is on extensive form games, that is dynamic games played out in time according to the extensive form, I will concentrate on the extensive form and forward induction refinements. Selten's subgame perfection is the first refinement in this category. The concepts that followed address games of incomplete information and are all designed to replicate the "incredible threat" clearing property of subgame perfection. Weak sequential equilibrium requires that there are beliefs that accord with strategies on reached information sets (but are otherwise unrestricted) and that actions must optimize payoffs given those beliefs. Perfect Bayesian equilibrium further tightens beliefs to accord with strategies if possible even on unreached information sets. The quest to tame those beliefs on unreached information sets continues next with sequential equilibrium [9], where beliefs are permitted if they survive in the limit of Nash equilibria of perturbed games with completely mixed strategies. The intuitive criterion [8] for signalling games restricts beliefs on off-equilibrium information sets by disallowing any weight on types that could not possibly profit³ from deviating. This was joined by many others including perfect sequential equilibrium [10], stable sets [8] and divine equilibrium [11].

2. Motivation and Examples

In the journey of refinements, perfecting belief formulation became an objective in itself, and perhaps, we lost sight of the reason why we need those beliefs.

Specifying off-equilibrium choices is certainly very useful for eliminating "incredible threats" and keeping strategies dynamically consistent.⁴ The conventional thinking was that in order to determine whether an action at an off-equilibrium information set is optimal, we need beliefs; and then, given those beliefs, the player making the choice can consider his/her expected payoff to make the optimal choice. Actions at unreached information sets that constitute an "incredible threat" are instances where a player is deliberately not maximizing at his/her decision point; the purpose being to thereby steer an earlier player away from that decision point for the threatening party's own advantage, when instead with the appropriate choice, the earlier player may have preferred to go to that decision point. Any other instances of non-optimal behaviour at unreached information sets does not constitute a threat. For the no threat information sets, non-optimal behaviour is not problematic and should not influence our understanding of what is a valid equilibrium. To illustrate this, some examples will be considered.

For simplicity, start with a full information game, the entry game in Figure 1, where $x \leq 2$. The entrant has two actions, I (opt In) or O (stay Out), and if he/she enters, the incumbent can either F (Fight) or A (Accommodate). Notice that there are two Nash equilibria, IA and OF. Fight (F) is an incredible threat,⁵ invalidating the OF equilibrium. We will call such an information set, or single decision point, where a player can make an incredible threat, pivotal. Now, if we alter the payoffs for

¹ A good literature survey on refinements can be found in Govindan and Wilson [5] or Sahanand and Sahanand [6,7].

² It is well understood that this is not a pure dichotomy. Generally, tremble-based refinements being normal form concepts mainly apply to static games, and forward induction refinements apply to dynamic games. However, there are exceptions. Extensive form trembling hand perfection though originating from the tremble philosophy exclusively applies to extensive form games. Iterated elimination of weakly-dominated strategies (which is not a true refinement, but rather a means to simplify a game), a completely normal form process, can at times coincide with forward induction solutions. Finally, the Kohlberg and Mertens stable sets [8] are also a normal form construct, but provide much insight into equilibria for incomplete information extensive form games.

³ Different refinements use different definitions of what is profitable. Differences are achieved by modifying the criteria of what constitutes the relevant set of best responses to types benefiting from choosing the out-of-equilibrium signal.

⁴ However, it is not the only way to eliminate incredible threats.

⁵ It is an incredible threat because if that point was actually reached, the incumbent would not fight, since his/her payoff is lower if he/she fights.

the challenger at Out (O) to $x \geq 2$, it becomes weakly dominant for the challenger to just staying out. It is not so important to specify the incumbent's behaviour because it is not relevant to the challenger's decision to stay out. Fight is no longer a threat that influences the challenger. The decision point is no longer pivotal. The incumbent's behaviour is not relevant to the challenger's decision to stay out. It would be fine to think that the incumbent plans to fight, or indeed not to think about the incumbent's choice at all, as in this game, it makes no difference.

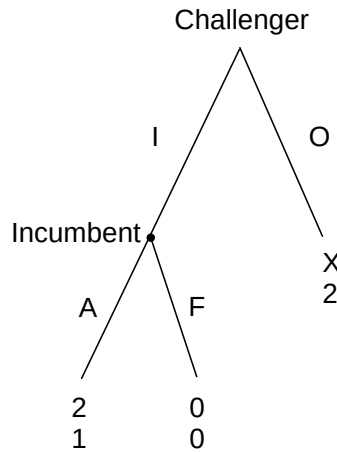


Figure 1. Entry game.

Consider the three-player game in Figure 2. Nash equilibria are Lab and Rbb. Here, we can agree that we should eliminate Rbb because Player 2 is using an incredible threat by saying that he/she plans to play b. By saying this, he/she is preventing Player 1 from choosing L, which he/she would have preferred. Therefore, in this instance, subgame perfection is important, and we do need Player 2 to make a rational choice at his/her decision point.




Figure 2. Three-player game with incredible threat.

Now, let us see a game without full information.⁶ Consider the game in Figure 3, with $x = 2$, $a = 1$ and $b = 0$.




Figure 3. Game with variable outside option.

There are two pure strategy equilibria here: (L, L') and (R, R') . If we include beliefs that rationalize Player 2's actions, we find the perfect Bayesian equilibria: $(L, L', p = 1)$ and $(R, R', p \leq \frac{1}{2})$. Furthermore, the second is problematic because the associated beliefs put much weight on action M, which is strictly dominated by R and weakly dominated by L, so it is not ever going to be played by Player 1. To solve this problem, a further restriction that is placed on beliefs is that there should be no probability weight put on the outcome of an action if it is strictly dominated at an earlier information set, if this is possible.⁷ This eliminates the (R, R') equilibrium, since now, Player 2 is not maximizing by choosing R' . While it is clear that the equilibrium needs eliminating, this seems like a complicated way to do it.

Consider some variations of this game. Suppose the payoff from the outside option R changed to $x \leq 1$. Now, action M is only weakly dominated $x = 1$, and if strictly less, it is not dominated at all. How can we eliminate the equilibrium $(R, R', p \leq \frac{1}{2})$, and do we want to eliminate it? A further modification to consider is that in addition to reducing the outside option payoff to $x = 1$, suppose that the payoffs after M for Player 1 are exchanged, so that $a = 0$ and $b = 1$. Does that change how we think about the (R, R') equilibrium? For all these versions of this game, with current refinements, we get different answers depending on the modification, but in each case, the only ideal reactive equilibrium is (L, L') .

Before we define ideal reactive equilibrium, consider a final example, given in Figure 4. Consider the two pure strategy Nash equilibria ALRa and ALLa. Conventional refinements would eliminate ALRa because Player 3's action is not rational; never mind that Player 3's action does not in any way constitute an incredible threat that influences the equilibrium path. Here is an instance where conventional theory requires too much rationality. In a sense, both are valid, because it does not matter what Player 3 plans to do. He/she could threaten by moving L, and it would be incredible; however, that is not what is motivating Player 1 to move A. With ideal reactive equilibrium, only the equilibrium path is relevant, because all the thinking that goes into the thought process dynamics has already distilled the incredible threats, and they are being ignored.

⁶ This is an example from Gibbons [12], p. 233.

⁷ "If possible" refers to the situation where both L and M were dominated by R, which would happen if we lower the payoff to Player 1 at L' to, say, $\frac{3}{2}$. Now, it is not possible to assign both nodes a probability of zero.




Figure 4. Incredible threat in this game poses no problem.

These are just a few examples. In what follows, we explore how to remove incredible threats without appealing to beliefs.

3. Ideal Reactive Equilibrium

Ideal reactive equilibrium utilizes the novel idea of the manipulated Nash thought process found in Amershi, Sadanand and Sadanand [1–4]. Although this work builds on that concept, it is distinct and involves much further development and investigation. Specifically, thought process dynamics itself is developed further here and in particular operationalized for signalling games. It is also analysed in greater detail from the perspective of first and second mover advantage games. Moreover, equilibrium is defined somewhat differently and shown to always exist for finite games.

To give a formal definition of ideal reactive equilibrium, we first need to define some terms that culminate in the definition of IRE.

1. Deliberate deviation of order n , $n \geq 1$: Starting from any set of initial choices for all the players, a deliberate deviation of order n by a particular player at a particular information set is a departure by that player from his/her initial choice to another of his/her available choices (pure action, no randomization), and he/she supposes that n of his/her immediate followers can observe the deviation. That is, we erase the information sets (if there are any) for those n players' choices that follow the particular player.

After a deliberate deviation, any player may respond with his/her best response to any changes that have occurred. This includes the original player. After that, any other player may respond with his/her best response to all the changes that have occurred thus far, and so on, until everyone is using a best response and we are at a (possibly different) Nash equilibrium.

2. Successful deliberate deviation: A deliberate deviation of any order is said to be successful if after the deviation and after all players are allowed to respond using best responses, the final outcome results in an increase in the payoff for the deviating player.

If a deliberate deviation is successful, then the outcome is that we remain at the final Nash equilibrium after deviation and best responses. If it is not successful, then the original deviator retracts his/her deviation, and we are back at the original starting point.

3. Thought process: A thought process is a sequence of outcomes starting with a deliberate deviation and then tracking all the best responses until we arrive at the resolution of that deviation. If the deviation is unsuccessful, then no movement occurs under the thought process; if it is successful, then the thought process tracks all the various stages of outcomes until we reach the final Nash equilibrium. Then, another deliberate deviation may be undertaken by any player, and the thought process continues to track the outcomes of that deviation and all the best responses that follow, as before. This continues until either there are no more successful deviations available or until we observe that the thought process is cycling indefinitely through a subset of the terminal nodes.

4. Ideal reactive equilibrium: A Nash equilibrium is an ideal reactive equilibrium if:

- all thought processes end with no change in the players' choices,
- or if all thought processes cycle endlessly through the same subset of terminal nodes, then the original Nash equilibrium is IRE if either its outcome is part of the subset through which the thought process cycles or, in the case of the original Nash equilibrium being a mixed strategy Nash equilibrium, if the cycling utilizes all the actions with positive probability weights in the mixed strategy.

As is evident from the definitions, the thought process gives importance to order of play, and earlier movers enjoy an advantageous position relative to later movers. Each player goes through successive orders of thinking⁸ and revising his/her choices, by imagining that arbitrary numbers of immediate followers can actually observe their moves even if in the original game, they could not. Upon "virtually observing",⁹ the followers may want to change their response.

This is all a thought process only. Players do not actually take any actions; it is only in the mind of the original player that began this line of thinking. Since all this thinking is simply based on the knowledge of the game tree, which all players possess, they can all anticipate each others' thinking. As indicated in the definitions, the original player may retract his/her deviation if he/she finds that he/she is actually worse off. Throughout the thought process dynamics, order of play is respected, and so by deviating, players can only influence their followers by supposing that they can observe the deviation, and not players preceding them. In case there are any nature moves, which are not strategic and cannot be influenced, players simply use expected payoffs in their thinking.

It may seem very strong to suppose that in a deliberate deviation of order n , the n follower players have a sudden and complete revelation of the deviator's new action. However, it should not be taken literally that the information is actually revealed, but rather, it is the deviator pondering: What if these n followers could figure out what I have done? Assuming they could figure it out, what would they then do? Then, what would everyone else do? At the end of it, if the deviator is indeed better off, then those n followers (who can also replicate the deviator's thought process) would expect the deviator to deviate, because he/she can and he/she is better off by doing it.

⁸ The process, though completely different, shares some similarity to the Wilson anticipatory equilibrium [13] for the Rothschild and Stiglitz incomplete information insurance model [14].

⁹ "Virtual observability" is a term coined by Camerer, Knez and Weber [15] in experiments where they tested the manipulated Nash thought process.

As we see from the definitions, beliefs are not necessary in defining and finding IRE. Payoffs depend only on the actions chosen by the players, as those are the only arguments of a payoff function. Probabilities can enter into the picture, and expected payoffs can be involved, but only to the extent that there are moves by nature. As we follow what happens to the original deviator, when we work through a thought process, as each stage, the actions chosen by everyone are well defined, and therefore, the payoffs for all players including the deviator can be computed, again without the need for any beliefs. Therefore, we can compute the deviator's payoffs due to the deviation and be certain whether he/she will find the deviation worthwhile. Similarly, we can compute payoffs at every stage of a thought process and finally determine which Nash equilibria are IRE, without ever appealing to beliefs.

Yet, we will see that in most instances, IRE achieves the same or a stronger refinement of Nash equilibrium, as compared to conventional belief-based refinements. The main intuition of how this is possible is the thought process dynamics. Recall that one of the features of any trembling hand refinement was to utilise completely mixed strategies. The idea was that every feasible action must have some probability weight so that the impact of that action on the whole game was not lost by focusing on some Nash equilibrium that put no weight on that action. Here, with IRE, instead of putting some weight on each action, the thought process dynamics allows each player to put all of his/her weight on each particular action by making a deliberate deviation towards that action. In doing so, the full effect of each possible action is considered, without requiring any beliefs.

A particular initial Nash action profile can be immune to deliberate deviation if all such forays into thinking eventually fail and all players retract their deliberate deviations to return to the original action profile. The other possible way that a deliberate deviation can be unsuccessful is if it does not lead anywhere because it keeps cycling. If every thought process fails because the deliberate deviations contained in it either return to the original action profile or cycle endlessly, then the original action profile becomes an ideal reactive equilibrium, because all deliberate deviations have failed to give the deviator a permanent improvement. Although an appealing feature of IRE is the simplicity of the thinking required and the similarity to how an uninitiated novice may approach the problem, it is very important to recognize that a fully-exhaustive search must be done over all possible thought processes. IRE is conceptually simpler, but since in most instances, it achieves a level of refinement that equals or exceeds conventional refinements, we should expect that this simplicity is balanced by the magnitude of the task of exhaustively searching through all possible thought processes.

To motivate the intuition of IRE, consider the following situation. Suppose two competitors discover two new markets where they can expand, market *A* and market *B*. They also learn that they each have resources to enter only one of the markets, and if both were in the same new market, they would compete away the profits and earn $(0, 0)$; however, if only one firm was in market *A*, it would earn three, and if only one firm was in market *B*, it would earn one. Firm 1 comes to know of this first and knows that he/she is first; and by the time Firm 2 learns, Firm 1 has already made secret and irreversible moves to enter one of the markets. The timing is very important, and the game tree is depicted in Figure 5.¹⁰

How should Firm 2 think about this? We know that if we solved for all the Nash equilibria, there are three: *AB*, *BA*, and a mixed strategy solution with each firm choosing *A* with probability 0.25. For *AB* to happen, Firm 1 thinks: Suppose I took the better market (*A*), would Firm 2 realize that that is what I would do, because being first, I would take the best thing available, and then figuring this out, Firm 2 would move to the other market (*B*)? Or should I take the worse market (*B*) in spite of being first, because when Firm 2 realizes all that has happened, he/she will think that I would leave the good market for him/her? Or should I randomize, and Firm 2 would be sure to figure that out and respond appropriately?

¹⁰ It is important to understand that a similar game tree with only the order of play reversed would be a completely different game that does not in any way depict our scenario, because the order of play there does not match the order of play given in the scenario.




Figure 5. New markets.

I would think it is quite obvious that the only thing Firm 2 can believe is that Firm 1 has taken the lucrative market, and so, it should make plans to enter the other market. This is essentially the battle of the sexes game. If we turn to the experimental evidence, Cooper, DeJong, Forsythe and Ross [16–18] found that when experimental subjects play the sequential unobserved battle of the sexes (without outside option, with outside option that is credible, with outside option that is not credible), they invariably choose the Player 1 preferred solution.¹¹ IRE predicts that Firm 1 would pick A and Firm 2 would pick B.

Accordingly, in the battle of the sexes game given in Figure 6, there are two pure strategy Nash equilibria, BB and SS, and a mixed strategy solution. The ideal reactive equilibrium, given that specific order of play, is only the BB solution. This is because SS is not immune to deliberate deviations, namely a first order deviation by Player 1 to B would be responded with B, and Player 1 would be better off, so he/she would undertake such a deviation. Instead, if we begin with the mixed strategy solution, again a first order deviation by Player 1 to B would be responded with B, and Player 1 would be better off, so he/she would undertake such a deviation. Therefore, both of these Nash equilibria are not IRE. Looking at the BB equilibrium, there is no deliberate deviation that any player can take that would make him/her better off. Player 2 cannot make any deliberate deviations, because he/she has no followers. Player 1 could deliberately deviate to S, but then, that would be met with a best response of S by Player 2, which would give Player 1 a payoff of only one, which is worse. Therefore, he/she would retract his/her deliberate deviation, confirming that BB is an IRE.

In effect, this equilibrium refinement extends the notion of forward induction to games where there is actually no signal, except the knowledge that earlier players move first and that confers an advantage to them. The order of play itself acts as a signal identifying which action a leader player must have taken.

It should be noted that the thought process does not always simply result in Player 1’s preferred equilibrium being selected. There are advantages in playing earlier in the game, namely that an early player can rule out large parts of the game by his/her choices. However, the disadvantage is that after moving, each player relinquishes control of the game to the players that remain. We see this in the three-player game in Figure 7.

¹¹ There are many other authors that have experimental results where they use other games (such as weak link games or resource dilemma games) and find that order of play does matter, such as Camerer, Knez and Weber [15], Budescu, Suleiman and Rapoport [19], Guth, Huck and Rapoport [20], and Muller and Sadanand [21]




Figure 6. Battle of the sexes.

There are two pure strategy Nash equilibria: LAa and RBb. In the spirit of virtual observability, it is almost as if Player 1’s move can be seen by the other players. We will argue that although LAa has a higher payoff for Player 1, RBb is the only ideal reactive equilibrium.

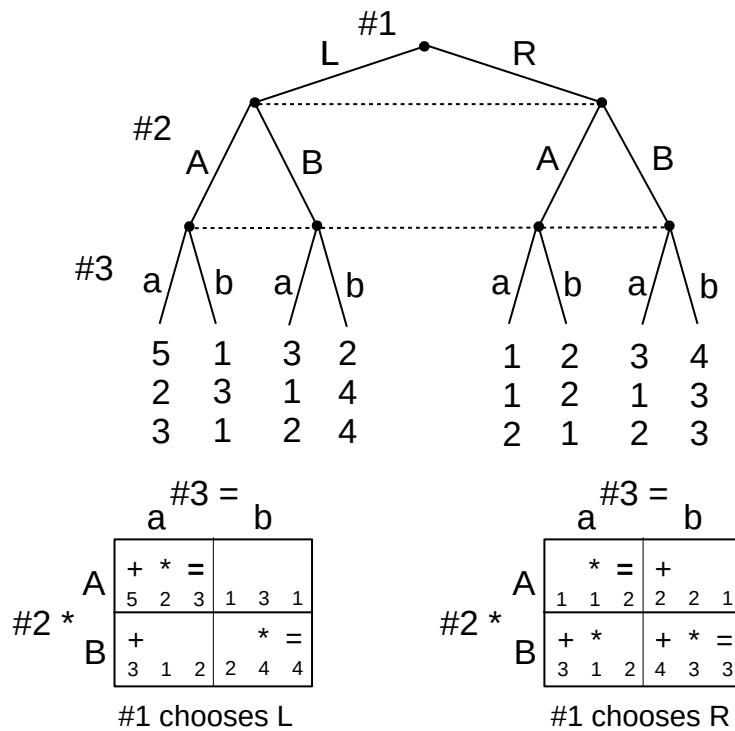


Figure 7. Three-player game.

Suppose we start at RBb and Player 1 considered a deviation to L. The players that follow would quite correctly anticipate it, as they should in the thought process dynamics. Given that Player 1 has already moved, he/she has committed to playing L and given up further control of the game. The other players would have anticipated Player 1’s move; therefore, they would behave as if they were playing the “subgame” after Player 1’s move L. By playing L, Player 1 has relinquished control of

the game, and Players 2 and 3 no longer need to worry about the Nash conditions that bind Player 1 to choosing L. He/she has already moved, it is finished; once moved, he/she cannot take it back. Therefore, now focusing on their “subgame”, Players 2 and 3 would figure out that it is in their best interest to play Bb, which is Nash in the “subgame.” This gives Player 1 a payoff of two, which is lower than his/her original payoff of four, under RBb. Player 1 can figure out all of this and would not make the deviation to L in the first place.

Exactly the same basic argument can be used to show that starting from LAa, there will be a successful deviation by Player 1 to R, and so, it is not an ideal reactive equilibrium.

The next example of ideal reactive equilibrium illustrates an instance where the thought process is not able to refine the Nash equilibria because the thought process dynamics in each case goes on endlessly without resolution, and so, the deliberate deviations are unsuccessful. Therefore, all of the Nash equilibria that can be reached in an endless thought process cycle are ideal reactive equilibria in this game.

Consider the game in Figure 8. There are two pure strategy Nash equilibria here: Lra and Rrb. However, any thought process dynamics from either starting point does not resolve. Lra brings us to terminal node z_3 , and from there, Player 2 could do an order one deviation to l, bringing us to z_2 . However, if that is what he/she is doing, Player 1 would make a best response bringing us to z_6 . Now, Player 3 would prefer to change, leaving us at z_5 , causing a change by Player 1 to z_1 , and then, Player 2 brings us back to z_3 , then the cycle could repeat again. This is of course not path independent, but as can easily be seen, all alternative paths also cycle endlessly, rendering all deliberate deviations unsuccessful. A similar result of unending cycling occurs for thought processes originating from Rrb, the other Nash equilibrium.




Figure 8. Thought process does not end.

One important test of whether an equilibrium concept is valid is whether real people play according to it. There is mounting experimental evidence that order of play makes a difference in sequential games where some or all of the moves by earlier players are not observed by follower

players. It is well established in the experimental literature that, contrary to conventional game theory, order of play affects how games are played. Normal form invariance is refuted.¹²

The first experimental results that suggested an order of play effect was the work by Cooper, DeJong, Forsythe and Ross [16–18]. There, they investigated the battle of the sexes game to see how experimental subjects dealt with outside options. They presented subjects with sequential unobserved battle of the sexes games without outside options, with outside options that were credible, and with outside options that were not credible. They expected to find that only in the games where the outside option was credible would you find a large propensity toward the Player 1 preferred equilibrium. They were puzzled to find that it did not matter whether the outside option was credible, or whether there even was an outside option. Players seemed to predominantly play the Player 1 preferred Nash equilibrium, which is exactly what IRE predicts. Camerer, Knez and Weber [15] set out to test whether real subjects played according to the Amershi, Sadanand and Sadanand [1–4] Mapnash equilibrium for weak link games such as the stag hunt. They found that in later rounds, their results confirmed a first mover advantage, which they called “virtual observability”. Muller and Sadanand [21] have found that subjects often choose in accordance with “virtual observability,” particularly for simple games. Brandts and Holt [24] and Rapoport [25] have also found order of play effects. In addition Schotter, Wegelt and Wilson [26] gave important results on how experimental subjects perceive order of play based on the presentation of the games.

The question that this growing strand of the experimental literature poses is: Exactly how does order of play affect outcomes? Indeed, Hammond [27] stated that the conventional assumption of normal form invariance may be “unduly restrictive” and emphasized the need for a systematic theory of how order of play affects outcomes.¹³ Huck and Müller [28] conducted some experiments and found that the opportunity to burn money confers a substantial advantage to first movers, and they too urge that “we make theoretical advances to understand the role of physical timing and first-mover advantages in games.” I believe that IRE is a good beginning at filling this void.

4. Properties and More Examples

There are four properties to notice about the ideal reactive equilibrium. First, the whole process happens without any need for beliefs. For unreached information sets that are not pivotal, it is not important to even specify behaviour for these counterfactual contingencies. Even when unreached information sets are pivotal, all players realize that they are so, and by working through all possible thought processes, they have already disallowed incredible threats, so again, it is not important to indicate choices there. Thus, the specification of what constitutes an equilibrium is simplified to almost the novice student’s approach of just providing the equilibrium path. Moreover, because sequential rationality-based refinements impose a blanket condition of rationality everywhere in the game, these conventional refinements can in some instances impose too much rationality and as a result make different predictions than what is predicted by IRE. The games where we see this are quite knife-edge, and for the most part except for such examples, IRE agrees with conventional belief-based refinements or further refines them. The example where we illustrate an instance of too much rationality is Selten’s horse, which will be discussed in detail in this section.

The second point is not a very precise one and will be illustrated in some of the examples to follow. The rough intuition is that if we can identify parts of games where preferences of players are aligned and where they are in conflict, then as long as the conflict payoffs are not higher, earlier

¹² Thompsen [22] and Dalkey [23] did some pioneering work establishing some fundamental invariances for extensive form games, which for years have been held as very basic tenets of game theory. Among them was normal form invariance. The idea was that games with the same normal form should have the same equilibrium solutions. However, experimental evidence shows that normal form invariance breaks down.

¹³ Hammond defines sophisticated Nash equilibrium as an answer to this question, but he states that “the precise relationship between sophisticated and manipulated Nash equilibrium deserves further exploration,” and thus, at this point, it is unknown.

players will, if possible, steer the game towards the parts where there is alignment, rather than conflict. The intuition for this is straightforward: the parts where there is conflict are second mover advantage situations where leader players hold no particular power and so can often result in randomization; the aligned portions are where leader players can exert some first mover advantage.

Next, because there exist knife-edge games where IRE and conventional refinements make completely different predictions, it will not be possible to show that one refines the other. However, except for the knife-edge games, in many instances, IRE actually refines belief-based refinements. Naturally, I will not be able to prove that IRE refines sequential rationality-based refinements, since there are examples where predictions differ. Instead, many examples, particularly of signalling games, will be provided to demonstrate that IRE refines belief-based equilibrium concepts. The examples of signalling games are mainly found in the next section.

Finally, I will show that IRE always exists for finite games. This is fairly easy to show by first noticing that since these are finite games, there is only a finite number of terminal nodes. Therefore, any thought process can only take a finite number of values, and therefore, they can only either cycle indefinitely or settle to a Nash equilibrium. From the definition, we know that if the thought process cycles indefinitely, there is an IRE. Therefore, we are left with the case where all thought processes settle to a Nash equilibrium. However, then, there must be an IRE at the Nash equilibrium where any one of them settles, because by definition, the settling point is where no more deliberate deviations are successful. If we started with such a Nash equilibrium, it would not be possible for any player to make any deliberate deviations that could be successful, and so, it would by definition be an IRE.

Returning now to the first property, consider Selten's horse in Figure 9. There are two equilibria DcL and CcR. Conventional theory argues that DcL should not be considered a valid outcome, because it requires Player 2 to choose c, which is not optimal when Player 3 is choosing L. However, the thing is that Player 2 is fully informed and knows perfectly well that he/she does not get a turn in the DL outcome, since Player 1 is playing D. Even Player 3 knows this and sets his/her beliefs correctly to $\text{prob}(x) = 1$ at his/her information set. If everyone were to use these correct beliefs, Player 2's action of playing c or d would be considered totally inconsequential. Moreover, we realize that Player 2's action c poses no "incredible threat" since switching to d (the "rational" choice) does not induce Player 1 to change his/her decision to exclude Player 2's decision node from the equilibrium path. This is because Player 1 would have worked out the entire thought process and realized that any change toward Player 2's decision node would trigger Player 3 to revise his/her choice from L to R, giving Player 1 a payoff of zero. Therefore, Player 1 would not entertain such a deviation in the first place.

Therefore, really, there is no need for Player 2's decision to be rational. Player 2's decision is unimportant, and this redeems the validity of DcL. Moreover, as it is the preferred outcome for Player 1, DcL is the equilibrium outcome selected by the ideal reactive equilibrium, which is contrary to most existing notions of equilibrium.¹⁴ This is an instance where conventional sequential rationality-based refinements impose too much rationality.

Consider the game of matching pennies in Figure 10, with an outside option paying (0.9,0.9), which being a second mover advantage game, illustrates the second point. The Nash equilibrium here is a mixed strategy equilibrium. Player 1 chooses O with certainty, and Player 2 mixes with any probability, $0.05 \leq \text{prob}(H) \leq 0.95$. In order to make these choices optimal for Player 2, the conventional approach is that we need beliefs on Player 2's information set that conform to Player 2's choice being optimal. We thus would have to assign the probabilities $\frac{1}{2}$ to each node. These beliefs, though quite plausible, are actually unnecessary, since if Player 2 held some other beliefs, that would result in no randomization, but as those beliefs themselves are unknown, Player 2's choice could still appear to be randomized (similar to Harsanyi's argument about mixed strategies). Player 1 is taking the outside

¹⁴ An interesting exception is the Gintis [29] local best response criterion, which agrees with ideal reactive equilibrium in its equilibrium selection for Selten's horse. However, it differs from the ideal reactive equilibrium in many ways, including its use of cardinal payoffs rather than ordinal ones, as seen from its treatment of the beer and quiche game.

option because it pays more, and so, he/she prefers to avoid the second mover advantage interaction with Player 2.

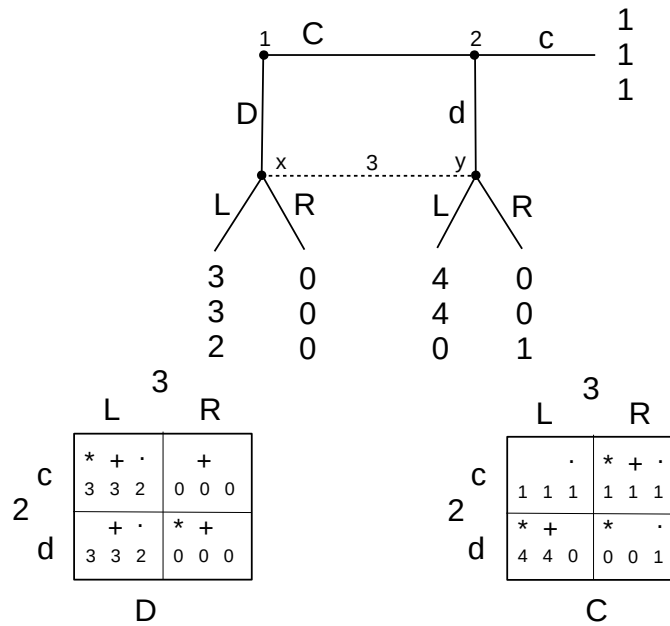


Figure 9. Selten's horse.




Figure 10. Matching pennies with outside option.

If instead, the outside option had paid the first player some negative amount, say -0.1 , then the game would play out differently. Now, although there is still an outside option, it is worse than what Player 1 expects to get from the second mover advantage interaction. The Nash equilibrium in this modified game will be that each randomizes with probability $1/2$ between H and T , and this is the IRE. To see that, consider all the possible deliberate deviations. Player 2 cannot make deliberate deviations since he/she has no followers, only Player 1 can. If Player 1 deviates to the outside option, he/she immediately retreats, since his/her expected payoff is lower. If instead, he/she deviates to H (or T), Player 2's best response will be T (or H), and this will be an endless cycle involving the actions over which they are randomizing, confirming that the original mixed strategy solution is IRE.

In games with a first mover advantage, it is beneficial to the first mover if the second mover could observe his/her move. We see that in many games where the first move cannot be completely observed, forward induction arguments often steer us to the first mover preferred outcome. If instead, it is a second mover advantage game, then forward induction is not very helpful in sorting out multiple equilibria, if there are any. Ideal reactive equilibrium also shares this tendency of other forward induction concepts.

The beer and quiche game [30], given in Figure 11, pairs both first and second mover advantage games with the two types of senders. We notice that the strong (\mathcal{S}) Player 1 plays a first mover advantage game (both with Beer (B) and with Quiche (Q), but prefers beer), while the weak (\mathcal{W}) Player 1 plays a second mover advantage game (again in both instances). In first mover advantage games, first players always want to reveal themselves as such, while in second mover advantage games, first players never want to reveal themselves. Therefore, it is not surprising that the chosen equilibrium is pooling, as the weak does not want to reveal himself, and the strong Player 1 gets his/her way with beer; if the weak were to try his/her way by choosing quiche, he/she would end up disclosing his/her identity. Then, the responder will clearly not duel for beer (N) and will duel (D) for quiche, which makes the weak sender worse off. The equilibrium response to both pooling with beer is no duel (N).




Figure 11. Beer and quiche.

We can go through the thought process dynamics to show that the Nash equilibrium where both senders pool at quiche is not an IRE. If we were in the pooling equilibrium at quiche, the strong sender would deviate to beer and will suppose that the receiver (follower) observes that he/she has done this. The receiver would then respond to beer by choosing N (no duel) and also realize that if anyone is signalling quiche, he/she must be the weak sender, and so also modify his/her response to quiche by D (duelling). Now, the weak sender will quickly change his/her choice to beer also. The thought process settles here because there are no more deliberate deviations that can be successful.

5. Signalling Games

In the previous section, we saw the beer and quiche game, which is a signalling game. We applied ideal reactive equilibrium informally there to show how we arrive at the same conclusion as the intuitive criterion, but in a much simpler fashion. In this section, we will discuss in more detail how to apply ideal reactive equilibrium to signalling games. We will see that in some instances, IRE can

refine further as compared to some standard sequential rationality-based signalling game refinements. For simplicity, we will suppose the game is a basic two-player game with the sender moving first and then the receiver.

To check whether a Nash equilibrium is ideal reactive, we begin a thought process where the sender of a particular type deviates deliberately and changes his/her message. Then, we follow all the possible best responses if the responder knew which type of sender it was. If any subgroup of senders is better off under the best responses for their type, then we continue the thought process. We also have to check whether the responder may change his/her reply to any of the other original signals being used in the Nash equilibrium, in light of this change. We also need to check whether any other sender types want to change their signal, and so on. This continues until all avenues have been explored. If none of the thought experiments lead to a change, the original Nash equilibrium is an ideal reactive equilibrium.

Consider the example in Figure 12. Each type is equally likely. The Nash equilibria are $mmmM$, $m'm'mL$, and $mm'mR$. From these, only $m'm'mL$ is an ideal reactive equilibrium.

		L	M	R
t_1	0,10	6,6	-1,5	-1,2
t_2	0,10	3,2	-1,1	1,2
t_3	0,10	-1,0	-2,3	-1,1
	m	m'		

Figure 12. Signalling game.

To see this, we will start with a Nash equilibrium that is not an IRE and go through the thought process to see how it arrives at the IRE. Suppose we start with $mmmM$, and suppose there is a deviation to m' possibly by t_2 . Now, let us look at the best responses for the receiver. $BR\{t_1\} = L$, $BR\{t_2\} = \{L, R\}$, $BR\{t_3\} = M$, $BR\{t_1, t_2\} = L$, $BR\{t_2, t_3\} = M$, and $BR\{t_1, t_3\} = M$. All players and types can immediately see that the deviator is not t_3 , since he/she is worse off with every response (equilibrium dominance). Therefore, it can only be t_1 or t_2 . Suppose that the responder thought it was just t_2 and the responder chose R as his/her response. Now, we are at $mm'mR$, which itself is a Nash equilibrium. However, the thought process does not end there. Type t_1 could do a first order deviation where he/she also deviates to m' , and now, the responder will be better off responding with L , rather than R . The thought process ends here, since no one wants to make a further deviation. Therefore, the original equilibrium $mmmM$ is not an ideal reactive equilibrium, nor is $mm'mR$. The only ideal reactive equilibrium is $m'm'mL$.

If we apply the Cho–Kreps intuitive criterion to this game, we can eliminate the pooling equilibrium $mmmM$. Notice that the beliefs supporting this pooling equilibrium at m' require that the receiver believes that the probability of $t_3 \geq 1/3$. However, equilibrium dominance indicates that this probability should be zero, which rules out the pooling equilibrium. However, beyond that, the Cho–Kreps intuitive criterion cannot refine any further, because it requires an off-equilibrium message in order to apply equilibrium dominance, and in the remaining two equilibria, all messages are used, so there is no out-of-equilibrium message. Similarly, divine equilibrium can also only eliminate the pooling equilibrium, because it too requires out of equilibrium messages to be able to assign sensible beliefs at unreached information sets. Thus, we see that IRE refines both the intuitive criterion and divine equilibrium.

Now, let us modify the example to allow the responder a move after m also, as seen in Figure 13. This example illustrates how it is not enough to simply follow what happens to the deviator and deviating message, but we must also see whether changes occur in the responder’s choices for types that do not deviate.

Here, the Nash equilibria are $mmmXM$ and $mm'mXR$, and only the latter is an ideal reactive equilibrium.

	X	Y	L	M	R
t_1	0,10	0,2	6,6	-1,5	-1,2
t_2	0,10	0,2	3,2	-1,1	1,2
t_3	0,2	-3,10	-1,0	-2,3	-1,1
	m		m'		

Figure 13. Modified signalling game.

Notice that the game after the message m' is identical to the last example. The only modification is that there is a choice between X and Y after the message m ; the payoffs under X are almost the same as before, and some new payoffs have been added for the response Y.

We want to show that the pooling equilibrium $mmmXM$ is not an ideal reactive equilibrium. We follow the same thought process for the previous example, with t_2 deviating to m' , and then t_1 , also trying to make that deviation, and the responder will choose L. It all seems good for the two deviating types.

However, notice now that if only t_3 is left signalling m , then the responder will change his/her response after m to Y.¹⁵ However, now, t_3 wants to deviate as well. Then, the responder’s best action after m' becomes M, which causes all of the types to retreat back to m . Since t_1 can anticipate all this, t_1 would not deviate to m' in the first place. Therefore, the thought process would rest at $mm'mXR$, which is the ideal reactive equilibrium. In this instance, IRE, the intuitive criterion and divine equilibrium all agree on $mm'mXR$.

These two examples serve to demonstrate how we conduct the thought process dynamics for signalling games. Essentially, it is the same requirement: a Nash equilibrium is ideal reactive if there is no successful deviation using thought process dynamics (or if all thought processes result in endless cycles). Only care must be taken to explore all possible avenues. They also illustrate how IRE compares with other conventional refinements.

6. Conclusions

This new equilibrium concept for dynamic games, ideal reactive equilibrium, returns to the basics after a long progression of refinements with increasing mathematical complexity. It gives importance to order of play, and it is almost a form of forward induction, as if earlier players can signal their

¹⁵ This is similar to what happened in the beer and quiche game. Starting with pooling at quiche, the strong switched to beer, and the receiver changed to no duel for beer. However, the receiver also switched his/her response to duel for quiche. Thus, it is important to also follow what changes happen to the receiver’s response to the original signal, as this induces further changes by the senders that were staying with the original signal. The changes that resulted in the beer and quiche example are that the other sender (weak) also joined the strong in messaging beer; much like in this example, t_3 also joined the others with the m' signal. However, for this example, it does not end there; a few more best responses occur before the thought process rests.

choices by virtue of their position in the order of play. Beliefs do not figure into this concept, and indeed, even specifying actions on the off-equilibrium path is unnecessary. In most instances, IRE seems to achieve very similar refinement, although in some instances, it is stricter and further refines conventional belief-based concepts, and in the occasional instance (in games where conventional concepts require too much rationality, such as Selten's horse), it is at odds with conventional results. While previous approaches to refinement were all stemming from the idea of extending subgame perfection to games with incomplete information and the need for reasonable beliefs to operationalize subgame perfection, here, with ideal reactive equilibrium, beliefs take a back seat. There is a new focus on thought process dynamics. The merits of this approach are the simplicity and the natural appeal of a process that is very aligned with how an uninitiated person would approach figuring out a game. Although seemingly very simple, this concept refines Nash equilibria, as well as many contemporary refinements and at times refines beyond conventional refinements. Finally, there is a growing body of experimental evidence that lends support to order of play effects in general, and also in particular to ideal reactive equilibrium. In reaction to the experimental literature, prominent theorists [27] have expressed that it is imperative to develop a theory of games that incorporates order of play, and IRE provides a good start at fulfilling that goal.

Funding: This research received no external funding.

Acknowledgments: I acknowledge the earlier work of Amershi, Sadanand and Sadanand (1985, 1992), from which this research begins, and makes a substantial extension and contribution. I also acknowledge assistance from Steven Gauthier in preparing the figures.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Amershi, A.; Sadanand, A.; Sadanand, V. *Manipulated Nash Equilibrium I: Forward Induction and thought Process Dynamics*; Faculty of Commerce and Business Administration University of British Columbia Working Paper 928; Faculty of Commerce and Business Administration, University of British Columbia: Vancouver, BC, Canada, 1985.
2. Amershi, A.; Sadanand, A.; Sadanand, V. *Manipulated Nash Equilibrium II: Properties*; Faculty of Commerce and Business Administration University of British Columbia Working Paper 1222a; Faculty of Commerce and Business Administration, University of British Columbia: Vancouver, BC, Canada, 1985.
3. Amershi, A.; Sadanand, A.; Sadanand, V. *Manipulated Nash Equilibrium III: Economic Applications and A Preliminary Experiment*; Faculty of Commerce and Business Administration, University of British Columbia Working Paper 1222b; Faculty of Commerce and Business Administration, University of British Columbia: Vancouver, BC, Canada, 1985.
4. Amershi, A.; Sadanand, A.; Sadanand, V. Player Importance and Forward Induction. *Econ. Lett.* **1992**, *38*, 291–297. [[CrossRef](#)]
5. Govindan, S.; Wilson, R. Refinements of Nash Equilibrium. In *The New Palgrave Dictionary of Economics*, 2nd ed.; Durlauf, S., Blume, L., Eds.; Palgrave Macmillan: London, UK, 2008; ISBN 978-0-333-78676-5.
6. Sadanand, A.; Sadanand, V. Equilibria in Non-Cooperative Games I: Perturbations Based Refinements of Nash Equilibrium. *Bull. Econ. Res.* **1994**, *46*, 197–224. [[CrossRef](#)]
7. Sadanand, A.; Sadanand, V. Equilibria in Non-Cooperative Games II: Deviations Based Refinements of Nash Equilibrium. *Bull. Econ. Res.* **1995**, *47*, 93–113. [[CrossRef](#)]
8. Kohlberg, E.; Mertens, J.-F. On the Strategic Stability of Equilibria. *Econometrica* **1986**, *54*, 1003–1038. [[CrossRef](#)]
9. Kreps, D.; Wilson, R. Sequential Equilibria. *Econometrica* **1982**, *50*, 864–894. [[CrossRef](#)]
10. Grossman, S.; Perry, M. Perfect Sequential Equilibrium. *J. Econ. Theory* **1986**, *39*, 97–119. [[CrossRef](#)]
11. Banks, J.; Sobel, J. Equilibrium Selection in Signaling Games. *Econometrica* **1987**, *55*, 647–661. [[CrossRef](#)]
12. Gibbons, R. *Game Theory for Applied Economists*; Princeton University Press: Princeton, NJ, USA, 1992; ISBN 978-0-691-00395-5.
13. Wilson, C. A model of insurance markets with incomplete information. *J. Econ. Theory* **1977**, *16*, 167–207. [[CrossRef](#)]

14. Rothschild, M.; Stiglitz, J. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Q. J. Econ.* **1976**, *90*, 629–649. [[CrossRef](#)]
15. Camerer, C.F.; Knez, M.; Weber, R. Virtual Observability in Ultimatum Bargaining and Weak Link Coordination Games. In Proceedings of the Meetings of the Public Choice Society and Economic Science Association, Houston, TX, USA, 12–14 April 1996.
16. Cooper, R.; DeJong, D.V.; Forsythe, R.; Ross, T.W. Selection Criteria in Coordination Games: Some Experimental Results. *Am. Econ. Rev.* **1990**, *80*, 218–233.
17. Cooper, R.; DeJong, D.V.; Forsythe, R.; Ross, T.W. Forward Induction in the Battle-of-the-Sexes Games. *Am. Econ. Rev.* **1993**, *83*, 1303–1315.
18. Cooper, R.; DeJong, D.V.; Forsythe, R.; Ross, T.W. Alternative Institutions for Resolving Coordination Problems: Experimental Evidence on Forward Induction and Preplay Communication. In *Problems of Coordination in Economic Activity*; Friedman, J.W., Ed.; Recent Economic Thought Series; Springer: Dordrecht, The Netherlands, 1994; Volume 35, Chapter 7, pp. 129–146.
19. Budescu, D.V.; Suleiman, R.; Rapoport, A. Positional and Group Size Effects in Resource Dilemmas with Uncertain Resources. *Organ. Behav. Hum. Decis.* **1995**, *61*, 225–238. [[CrossRef](#)]
20. Guth, W.; Huck, S.; Rapoport, A. The Limitations of the Positional Order Effect: Can it Support Silent Threats and Non-Equilibrium Behavior? *J. Econ. Behav. Organ.* **1998**, *34*, 313–325. [[CrossRef](#)]
21. Muller, A.; Sadanand, A. Order of Play, Forward Induction, and Presentation Effects in Two-Person Games. *Exper. Econ.* **2003**, *6*, 5–25. [[CrossRef](#)]
22. Thompson, F.B. Equivalence of Games in Extensive Form. *Rand Res. Memo.* **1952**, *RM-759*, 1–13.
23. Dalkey, N. Equivalence of Information Patterns and Essentially Determinate Games. In *Contributions to the Theory of Games*; Kuhn, H., Tucker, A., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; Volume 2, pp. 217–244.
24. Brandts, J.; Holt, C.A. Forward Induction: Experimental Evidence from Two-Stage Games with Complete Information. In *Research in Experimental Economics*; Isaac, R.M., Ed.; Emerald Publishing Limited: Bingley, UK, 1992; Volume 5, 119–137.
25. Rapoport, A. Order of Play in Strategically Equivalent Games in Extensive Form. *Int. J. Game Theory* **1997**, *26*, 113–136. [[CrossRef](#)]
26. Schotter, A.; Weigelt, K.; Wilson, C. A Laboratory Investigation of Multiperson Rationality and Presentation Effects. *Games Econ. Behav.* **1994**, *6*, 445–468. [[CrossRef](#)]
27. Hammond, P. *Beyond Normal Form Invariance: First Mover Advantage in Two-Stage Games with or without Predictable Cheap Talk*; Economic Research Papers 269784; Department of Economics, University of Warwick: Coventry, UK, 2008.
28. Huck, S.; Müller, W. Burning Money and (Pseudo) First-Mover Advantages: An Experimental Study on Forward Induction. *Games Econ. Behav.* **2005**, *51*, 109–127. [[CrossRef](#)]
29. Gintis, H. The local best response criterion: An epistemic approach to equilibrium refinement. *JEBO* **2009**, *71*, 89–97. [[CrossRef](#)]
30. Cho, I.-K.; Kreps, D. Signalling Games and Stable Equilibria. *Q. J. Econ.* **1987**, *102*, 179–222. [[CrossRef](#)]



© 2019 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).