

UC Irvine

UC Irvine Previously Published Works

Title

Generalized Backward Induction: Justification for a Folk Algorithm

Permalink

<https://escholarship.org/uc/item/1wm267w3>

Journal

Games, 10(3)

ISSN

2073-4336

Author

Kaminski, Marek Mikolaj

Publication Date

2019

DOI

10.3390/g10030034

Peer reviewed

Article

Generalized Backward Induction: Justification for a Folk Algorithm

Marek Mikolaj Kaminski 

Department of Political Science and Mathematical Behavioral Sciences, University of California, 3151 Social Science Plaza, Irvine, CA 92697-5100, USA; marek.kaminski@uci.edu

Received: 11 June 2019; Accepted: 27 August 2019; Published: 30 August 2019



Abstract: I introduce axiomatically infinite sequential games that extend Kuhn’s classical framework. Infinite games allow for (a) imperfect information, (b) an infinite horizon, and (c) infinite action sets. A generalized backward induction (GBI) procedure is defined for all such games over the roots of subgames. A strategy profile that survives backward pruning is called a backward induction solution (BIS). The main result of this paper finds that, similar to finite games of perfect information, the sets of BIS and subgame perfect equilibria (SPE) coincide for both pure strategies and for behavioral strategies that satisfy the conditions of finite support and finite crossing. Additionally, I discuss five examples of well-known games and political economy models that can be solved with GBI but not classic backward induction (BI). The contributions of this paper include (a) the axiomatization of a class of infinite games, (b) the extension of backward induction to infinite games, and (c) the proof that BIS and SPEs are identical for infinite games.

Keywords: subgame perfect equilibrium; backward induction; refinement; axiomatic game theory; agenda setter; imperfect information; political economy

JEL Classification: C73

1. Introduction

The origins of backward induction are murky. Zermelo [1] analyzed winning in chess, asking a question about the winning strategy for white in a limited number of moves, yet his method of analysis was based on a different principle [2]. A couple of decades later, we find that reasoning based on backward induction was implicit in Stackelberg’s [3] construction of his alternative to the Cournot equilibrium. Then, as a general procedure for solving two-person, zero-sum games of perfect information, backward induction appeared in von Neumann and Morgenstern’s [4] (p. 117) founding book, which followed von Neumann’s [5] earlier question about optimal strategies. BI was also used to prove a precursor of Kuhn’s Theorem for chess and similar games. Von Neumann’s exceedingly complex formulation was later clarified and elevated to high theoretical status by Kuhn’s [6] work, most especially, Corollary 1; Schelling’s [7] ideas about incredible threats; and Selten’s [8] introduction of subgame perfection. However, it suffered drawbacks when the chain-store paradox, Centipede, and other games brought into doubt its universal appeal [9,10]. Its profile was further lowered with new refinements: perfect equilibrium [11], sequential equilibrium [12], and particular procedures such as forward induction [13] offered solutions that often seemed more intuitive. The arguments against applying backward induction began to multiply, sealing doubts about its universal validity [14–22].

The goal of the present paper is to extend backward induction to infinite games with imperfect information and to investigate its relation to subgame perfection. In its standard formulation, backward induction applies only to finite games of perfect information. In these cases, every backward induction solution (BIS), i.e., a strategy profile that survives backward “pruning” (a subsequent substitution of

terminal subgames with Nash-equilibrium payoffs), is also a subgame perfect equilibrium (SPE) and all SPEs result from backward pruning [23] yet it is common knowledge among game theorists that other games can be solved through backward reasoning as well; they have routinely applied the procedure to such games for half a century. Backward reasoning is implicit in refining the Stackelberg equilibrium from other Nash equilibria (NE). Schelling analyzed backward the NE in the iterated Prisoner’s Dilemma as early as in the 1950s [24]. Reputable textbooks, such as those by Fudenberg and Tirole [25] (p. 72) and Myerson [26] (p. 192), make explicit claims—although without proofs—that backward induction can be applied to a wider class of games. Osborne and Rubinstein [27] (p. 98, Lemma 98.2 on one deviation property) essentially extend backward induction to finite-horizon extensive games with perfect information. Additionally, Escardó and Oliva [28] showed that BIS is equivalent to SPE in certain well-founded games, i.e., games with perfect information, such that all paths must be finite, even though they may be arbitrarily long. Moreover, backward reasoning is commonly applied to solve various parametrized families of extensive-form games in political economy models. It is also used implicitly when one argues that voters “vote sincerely” in the last stage of various voting games that have binary agendas.

What results from this alleged abuse? Perhaps most clearly, Fudenberg and Tirole [25] (p. 94) spell out the underlying principle: “This is the logic of subgame perfection: Replace any ‘proper subgame’ of the tree with one of its Nash-equilibrium payoffs, and perform backward induction on the reduced tree.” With caveats discussed later in this introduction, Fudenberg and Tirole’s prescription captures the essence of the generalized backward induction (GBI) approach. However, if one wanted to find an entirely formal (i.e., axiomatic) justification for this algorithm in the literature, one would not be able to do so, yet such a justification is necessary since, in fact, the principles behind SPE and BIS differ.

We can informally examine this *petite difference*. SPE is a strategy profile that is NE in all subgames. Its intuitive justification focuses on subgames, i.e., smaller games within a larger one. SPE demands that players interact “rationally” in all subgames, i.e., that they apply strategies resulting in NE in the subgames. Unless a game has no other subgames than itself, backward induction concerns different games. It starts at the game’s end and moves backward. Similar to SPE, the first game (or a set of games) is a subgame. However, at some point, a new game appears that is not a subgame of the original game and which has no counterpart in the SPE’s definition. Such a game (or games) is created by substituting a subgame (or subgames) with an NE payoff vector in those subgames.

A simple example illustrates the distinction. In Figure 1, SPE in game G requires that Bob and Alice play NE strategies in both the original game G and its subgame H. BIS requires that the players play NE strategies in G’s subgame H and also the “upgame” J. J was created by “pruning” H, i.e., substituting the root of H with an endnode that had the NE payoffs in H assigned as payoffs.




Figure 1. Backward induction versus subgame perfection.

Backward induction is helpful because the games that we need to solve when using its algorithm are usually smaller and simpler than those that appear in the definition of the SPE. It seems such an obvious

route toward finding SPEs that one can easily ignore the subtle difference in the definitions, especially when one begins to struggle with the truly painful chore of defining BIS formally. Nevertheless, one cannot state a priori that the sets of strategy profiles obtained in both cases are identical.

Despite the lack of a firm foundation, it seems that, within the discipline, backward induction has acquired the status of a “folk algorithm.” Although game theorists use it, it appears that little attention has been paid to its rigorous examination. It clearly works—in the sense that it produces SPEs—but we neither know exactly for what games it works nor precisely what it produces. The formal link between the folk algorithm and subgame perfection remains obscure.

In the current paper, I investigate the “folk puzzle of backward induction” axiomatically. The axiomatic framework that I introduce below encompasses more games than the finite games axiomatized by either von Neumann and Morgenstern [4] or Kuhn [6]. Specifically, in this work, the action sets may be infinite (more formally, they can have any cardinality) and the length of a game may be infinite (denumerable). Despite its higher complexity, this new framework allows for a constructive analysis.

The results of this work certify that, with some clarification, the folk wisdom linking BI and SPE is correct for various sets of infinite games. Nevertheless, it is also clear that Fudenberg and Tirole’s [25] prescription of applying BI over subgames has to be modified. The informal description of these modifications is as follows:

First, backward pruning can be applied not only to pure strategies but also to certain behavioral strategies.

Second, we must replace every subgame with an SPE and not an NE payoff vector. Obviously, when a subgame has no other subgames than itself, SPE and NE coincide.

Finally, we can replace entire subsets of disjoint subgames simultaneously, not only single subgames. Such simultaneous replacement is essential not only for solving many games that are not finite but also for certain complex finite games, as it may be the only realistic way to proceed (see Example 1).

Remarkably, in agreement with the case of finite games of perfect information, the findings show that, for pure strategies and certain behavioral strategies, the sets of BIS and SPE coincide. This is the main result of the present work. It legitimizes the informal methods of backward pruning of a game, and it concatenates the resulting partial strategy profiles used by game theorists to solve games.

The generalized algorithm uses the agenda—the tree consisting of the roots of all subgames—instead of the game tree. For games of perfect information, the agenda coincides with the game tree with terminal nodes subtracted. A step in such an algorithm can informally be compared to the classic backward induction as follows:

- (1) Prune (remove) any subset of disjoint subgames instead of a single subgame, which would have only one decision node followed by terminal nodes.
- (2) Substitute all selected subgames with the SPE payoffs instead of the payoffs for best moves.
- (3) Concatenate all partial strategy profiles obtained in the previous step; if at any point one gets an empty set, this would mean that there is no SPE in the game.

The above three-step procedure can be applied to pure strategies in all games as well as to behavioral strategies that satisfy the condition of finite support and crossing. Finding all SPEs requires following particular rules for concatenating and discarding partial strategy profiles, as described in Section 4 of this paper.

The next section introduces axiomatically sequential games with potentially infinite sets of actions and an infinite horizon. While a more accurate label would be “potentially infinite” here, I call such games “infinite” for the sake of simplicity. Following the introduction of these games, I establish the basic facts linking payoffs to strategies in infinite games. Then, in Section 3, I investigate the decomposition of games into subgames and upgames for pure and behavioral strategies with finite support and crossing. The reader who is more interested in applications may skip Sections 2 and 3

and go directly to Section 4, where the GBI algorithm is formally described. Then, in the subsequent section, I discuss five illustrative examples, including the application of the procedure to solve parametrized games, which are often used in political economy modeling and an ad hoc generalization that finds the unique perfect equilibrium. Finally, I conclude the paper in Section 6 with the remaining open questions as well as suggestions for further research. All proofs are found in the Appendix A.

2. Preliminaries

Sequential (extensive-form) games were introduced with a set-theoretic axiomatization by von Neumann and Morgenstern [4] (pp. 73–76). They were conceived and presented in the spirit of the rigorous decoupling of syntax and semantics in the early 20th century, such as embodied in the works of Hilbert and Tarski and, later, in Bourbaki's team. In order to prevent a reader from forming any geometric or other intuition, von Neumann [4] proudly announced that “we have even avoided giving names to the mathematical concepts [...] in order to establish no correlation with any meaning which the verbal associations of names may suggest” (p. 74). Then, he dismissed his own idea of a game tree, as “even relatively simple games lead to complicated and confusing diagrams, and so the usual advantages of graphical representation do not obtain” (p. 77). Despite von Neumann's every effort to turn a sequential game into a highly abstract object, incomprehensible for nonmathematicians, Kuhn's [6] approach helped to make games more intuitive. Kuhn simplified von Neumann's formalism and built the axioms into definitions and assumptions about the tree, the players, and the information, as well as slightly generalized von Neumann's unnecessarily narrow definition.

The axiomatic setup of the current paper goes beyond finite games in an attempt to cover axiomatically a larger class of extensive-form games. As mentioned above, for simplicity, such games are called *infinite games*. The framework maintains the compatibility with Kuhn's pragmatic exposition and draws from the excellent modern presentations of finite games by Myerson [26] and Selten [11]. The axioms are divided into two subsets. The properties defining an infinite tree are listed explicitly; the game axioms are combined with the description of the game components and specify how various objects are attached to the tree. In order to establish some intuitive associations, I place in parentheses the axioms' names that succinctly describe their content.

The way in which games are introduced is laborious, but it helps later with the succinct establishment of the basic results.

Rooted tree: Let (T, Ψ, τ) be such that T is a set of at least two points, Ψ is a binary relation over T , and $\tau \in T$. For $y \in T, y \neq \tau$, a *path* to y of length k is any finite indexed set $e_y = \{x_i\}_{i=1}^k \subset T$ such that $x_1 = \tau, x_k = y$ and for all $i = 1, \dots, k-1, (x_i, x_{i+1}) \in \Psi$; for $i \neq j, x_i \neq x_j$. For $y = \tau$, the path to τ is $\{\tau\}$. For an infinite indexed set, $\{x_i\}$ is an infinite path if for every $x_j \in \{x_i\}, \{x_i\}_{i=1}^j$ is a finite path.

Ψ is termed a *rooted tree* with its *root* τ and the set of *nodes* T if the following properties (AT1-AT3) are satisfied:

AT1 (*partial anti-reflexivity*): For all $x \in T, (x, x) \in \Psi$ iff $x = \tau$;

AT2 (*symmetry*): For all $x, y \in T (x, y) \in \Psi$ iff $(y, x) \in \Psi$;

AT3 (*unique path*): For every $x \in T$, there is exactly one *path* e_x to x .

The following definitions and notation are used hereafter (the definitions are slightly redundant):

1. Binary relations between two different nodes x, y :

Predecessor: $y \in PR(x)$ (*precedes* x or *is in the path to* x) iff $e_y \subset e_x$;

Successor: $y \in SU(x)$ (*follows* x) iff $e_x \subset e_y$;

Immediate predecessor: $y = IP(x)$ (*immediately precedes* x) iff $e_x - e_y = \{x\}$

By AT1-AT3, for every $x \neq \tau$, there is exactly one immediate predecessor.

Immediate successor: $y \in IS(x)$ (*immediately follows*) x iff $e_y - e_x = \{y\}$;

Immediate predecessor in $T_i \subset T$: For $x, y \in T_i, y = IP_i(x)$ (*immediately precedes* x in T_i) iff $y \in PR(x)$ and $(e_x - e_y) \cap T_i = \{x\}$.

By AT1-AT3, for every x and T_i , there is at most one immediate predecessor of x in T_i . This definition allows us to import the relation of preceding from the game tree into a smaller tree consisting of the roots of the game's subgames (i.e., the game's *agenda*). GBI will be conducted over the agenda.

2. Single nodes, subsets of nodes, set of subsets of nodes

Endnode (terminal node): a node x that is not followed by any other node, i.e., $SU(x) = \emptyset$;

Set of all endnodes: T_E ;

Decision node: a node that is not an endnode;

Set of all decision nodes: $T_D = T - T_E$;

Branch: any node except for the root τ . (In a standard definition, a branch is any element $(x, y) \in \Psi$. For the simplicity of forthcoming notation, it is identified here with the node in the pair that is farther from the root.)

Alternative (originating) at a node x : any immediate successor of x ;

Terminal path: a path to an endnode or an infinite path;

Set of all terminal paths T_t .

Definition 1. *Game*: An n -player sequential game is a septuple $G = \langle \Psi, N^0, \{T_i\}_{i \in N^0}, I, A, h, P \rangle$ that includes a rooted game tree Ψ and the following objects: players from N^0 with their assigned decision nodes and probability distributions for random moves $\{T_i\}_{i \in N^0}$; the pattern of information I ; the identification of moves A ; the probability distributions over random (or pseudorandom) moves h ; and the payoff function P .

The conditions imposed on the components of G and certain useful derived concepts are defined below.

1. *Game tree*: Ψ is a rooted tree with a set of nodes T , a set of decision nodes T_D , a set of endnodes T_E , and the root τ ;
2. *Players*: For a positive integer n , $N^0 = \{0, 1, \dots, n\}$ consists of *players* $N = \{1, \dots, n\}$ and nature—a *random* or *pseudorandom mechanism*, labeled with 0;
3. *Player partition*: $\{T_i\}_{i \in N^0}$ is the partition of T_D into (possibly empty) subsets T_i and a (possibly empty) subset T_0 for nature. The following assumptions are made regarding T_0 :

- (i) There is no path that includes an infinite number of nodes from T_0 ; 1'
- (ii) For every $x \in T_0$, the number of alternatives at x is greater or equal two and finite.

4. *Information*: $I = \{I_i\}_{i=0}^n$ is such that every $I_i = \{I_i^k\}_{k \in K_i}$ is a partition of i 's set T_i . We assume the following:

- (i) All elements of I_0 are singletons;
- (ii) For all $i \in N$, every element of I_i includes only nodes with equal numbers (or cardinalities) of alternatives and does not include two nodes that are in the same path;
- (iii) (*Perfect Recall*): If x_i is a successor of x_j and x_k is in the same information set with x_i , then x_i and x_k must be immediate successors of either x_j or some other node x_l that is a successor of x_j .

For every $i \in N^0$, a set $I_i^k \in I_i$ is called i 's *information set*. A node y *originates* from the information set I_i^k if y originates at a node $x \in I_i^k$.

5. *Moves (actions)*: $A = \{A_i^k\}_{i \in N, k \in K_i}$ is a collection of partitions, one for every information set I_i^k of every player i , of all alternatives originating at I_i^k , such that for any node $x \in I_i^k$, every member of A_i^k includes exactly one alternative that originates at x . The elements $a \in A_i^k$ are called the *moves (or actions)* of i at I_i^k . For any I_0^k , the moves at I_0^k are singletons, including branches originating at

- I_0^k . By definition, since every branch y belongs to precisely one move, for $y \in T - \{\tau\}$ the move a such that $y \in a$ is denoted by $A_i^k(y)$.
6. *Random moves:* h is a function that assigns to every information set of the random mechanism $I_0^k = \{x\}$ a probability distribution $\{h^k\}$ over the alternatives at x , with all probabilities being positive. If $T_0 = \emptyset$, h is not defined.
 7. *Payoffs (associated with terminal paths):* The *payoff function* $P = (P_1, \dots, P_n) : N \times T_t \rightarrow \mathbb{R}$ assigns to every terminal path $e \in T_t$ a *payoff vector* at e equal to $P(e) = (P_1(e), \dots, P_n(e))$. The component $P_i(e)$ is called the *payoff* of player i at e . Function P_i is called the *payoff function* of player i .

Both infinite paths and infinite numbers of moves at the players' information sets are allowed. The assumed constraints demand that the numbers of players, the random information sets at every path, and the random moves at every random information set are finite. Another restriction is the discrete temporal structure of moves implied by the definition of game tree. Such a restriction excludes differential games and, in general, games in continuous time. Finally, games like stochastic games are conceptualized differently than extensive-form games (see, e.g., Reference [29]). Hereafter, the word "game" refers to Definition 1.

The most extensively studied subset of infinite games are finite games:

Finite game: G is finite if its set of nodes T is finite.

The concepts that follow are derived from the model's primitives.

Subgame: For any game $G = \langle \Psi, N^0, \{T_i\}_{i \in N^0}, I, A, h, P \rangle$, a subgame of G is any game $G' = \langle \Psi', N^0, \{T'_i\}_{i \in N^0}, I', A', h', P' \rangle$, such that

- (i) Ψ' is a subtree of Ψ , i.e., for some $\tau' \in T$, $T' = \{x \in T : x = \tau' \text{ or } x \in SU(\tau')\}$ and $\Psi' = (\Psi \cap [T' \times T']) \cup \{(\tau', \tau')\}$;
- (ii) If $x_1, x_2 \in I_i^k$ for some I_i^k in G , then either $\{x_1, x_2\} \subset T'$ or $\{x_1, x_2\} \cap T' = \emptyset$ (either both x_1 and x_2 are in T' or neither is);
- (iii) The sets of players are identical (N^0) and $\{T'_i\}_{i \in N^0}, I', A', h', P'$ are restrictions of $\{T_i\}_{i \in N^0}, I, A, h, P$ to T' , respectively.

It is straightforward that restrictions in (iii) define a game and that "being a subgame" is a transitive relation.

A player i may be a dummy in a game, i.e., T_i may be empty. Since $|T| \geq 2$, the root of Ψ is a decision node and there must be at least one player or random mechanism in the game. Without loss of generality, one can assume that there are no dummies in the initial game G for which all results are formulated.

Below, the strategies are defined in order to optimize the introduction of the fundamental ideas—for this paper—of strategy concatenation and decomposition. The adjective "behavioral" is optional since behavioral strategies are our departure point for defining other types of strategies.

Behavioral actions: A *behavioral action of finite support* (in short, a *behavioral action*) α_i^k of player i at the information set I_i^k is a finite probability distribution over the set of actions A_i^k .

A behavioral action at any information set assigns positive probabilities only to a finite number of actions at this set. When I refer to "finite support," I will also mean the assumption in the definition of a game that the number of random actions is also finite.

Strategies (rough behavioral): A *rough behavioral strategy* β_i of player i is any (possibly empty) set of i 's behavioral actions that includes exactly one behavioral action per information set of i . A *partial rough behavioral strategy* ω_i is any subset of a rough strategy. A partial rough strategy that includes exactly those actions in β_i that are defined for information sets of i in a subgame H of G is denoted as β_i^H and is called β_i reduced to H .

The possibility of having an empty set of i 's behavioral actions represents a trivial strategy of player i in a subgame where i makes no moves. A rough behavioral strategy may be assembled from any set of actions. Below, this option is restricted to rough strategies that satisfy finite crossing.

For any rough strategy β_i , we denote the probability assigned by β_i at I_i^k to a move a_i by $\beta_i(I_i^k)(a_i)$. A path e is called *relevant* for β_i if β_i chooses every alternative in e that originates from some information set of i with a positive probability, i.e., if for every node $y \in e$ such that $y \in IS(x)$ for some $x \in T_i$, $\beta_i(I_i^k)(A_i^k(y)) > 0$. Finally, a path e *crosses* I_i^k if $e \cap I_i^k \neq \emptyset$.

Finite crossing in subgames: For every $i \in N$, a rough strategy β_i is said to satisfy finite crossing in the subgames if for every subgame H of G and every path e^H in H that is relevant for β_i reduced to H , β_i^H, e^H crosses only a finite number of information sets from I_i , such that the distribution $\beta_i(I_i^k)$ is nondegenerate, i.e., it assigns positive probabilities to at least two actions.

Behavioral strategy: For any $i \in N$, a behavioral strategy of i —or simply, a strategy of i —is any rough strategy of i that satisfies *finite crossing in subgames*.

Comment: Finite support and finite crossing guarantee that, in all subgames, the payoffs (to be defined later) for behavioral strategies can be derived from a finite probability distribution. The class of strategies that satisfy these conditions includes, among others, all behavioral strategies in a finite game and all pure strategies in any game. Relaxing these conditions would introduce complications of a measure-theoretic nature along the lines examined by Aumann [30]. It is unclear whether the results would survive a more general treatment of behavioral strategies.

All behavioral strategies of i form i 's *behavioral strategy space* B_i . Elements of $B = \times_{i=1}^n B_i$, are called *behavioral strategy profiles* and are denoted by β .

I use a set-theoretic interpretation of strategies that will greatly simplify the definitions of strategy decomposition and concatenation as well as the treatment of partial strategy profiles. There is a simple isomorphism between strategy profiles defined in a set-theoretic and standard way. Thus, every strategy profile β is interpreted as a union of players' strategies (which are obviously disjoint); the Cartesian product $\times_{i=1}^n B_i$ is interpreted as taking all possible unions of individual strategies, one per player; the notation for the strategy profile $(\beta_i)_{i=1}^n$ represents an alternative notation for $\cup_{i=1}^n \beta_i$. One example of the notational difficulty that is avoided is the interpretation of $\times_{i=1}^n B_i$ when at least one strategy set is empty. Another example is provided by the next definition.

A strategy profile β with the strategy of player i removed, i.e., $\beta - \beta_i$, is denoted by $\beta_{-i} \in \times_{N - \{i\}} B_i$; (β_{-i}, γ_i) denotes β with β_i substituted with γ_i , that is, $\beta - \beta_i \cup \gamma_i$.

When such a distinction is necessary, the payoff functions, strategies, strategy profiles, etc. in games or in subgames G and H will be given the identifying superscripts P^G, P^H , etc.

The most important step toward building the framework for infinite games is expressing payoffs in terms of strategies.

Recall that the probability assigned by β_i at I_i^k to a move a_i was denoted by $\beta_i(I_i^k)(a_i)$. For any $x, y \in T$ —such that $x \in IS(y)$ and $y \in I_i^k$ —and $x \in a_i$, the probability of *the move to* x is defined as $p_\beta^m(x) = \beta_i(I_i^k)(a_i)$. By convention, $p_\beta^m(\tau) = 1$ for all β . A path e is *included* in β if $p_\beta^m(x) > 0$ for all $x \in e$, and it is denoted as $e \subset \beta$. The set of all terminal paths included in β is denoted by T_β . The probability of *playing* e under β , $p_\beta(e)$, is defined as follows:

$$p_\beta(e) = \prod_{x \in e} p_\beta^m(x)$$

Thus, $p_\beta(e)$ is the product of the probabilities assigned by β to all alternatives in e . Note that by the definition of the game and the behavioral strategy's finite support of crossing, for a path of infinite length, only a finite number of alternatives may be assigned probabilities different than zero or one. The multiplication over an infinite series of numbers has at least one zero equal to zero, and an infinite product of ones is equal to one.

The probability of *reaching* a node y , $p_\beta(y)$, is defined as the probability of playing e_y under β :

$$p_\beta(y) = p_\beta(e_y) = \prod_{x \in e_y} p_\beta^m(x).$$

The assumptions of finite support and finite crossing are used below to establish the fundamental fact that p_β defines a probability distribution over a finite subset of all terminal paths.

Lemma 1. For every game G and every subgame H of G and with the behavioral strategy profile β ,

- (a) the set of all terminal paths in H included in β^H, T_{β^H} , is non-empty and finite;
- (b) for every $e^H \in T_i^H, p_\beta^H(e^H) > 0$ iff $e^H \in T_{\beta^H}$;
- (c) $\sum_{e \in T_{\beta^H}} p_{\beta^H}(e) = 1$.

Lemma 1 establishes that the probability distributions associated with actions of every behavioral strategy in a profile define a finite probability distribution on the set of all terminal paths. This allows us to extend the definition of the payoffs that were originally defined only for terminal paths to all behavioral strategy profiles of finite support and finite crossing.

Payoffs (for behavioral strategy profiles): For every behavioral strategy profile $\beta \in B, P_i(\beta) = \sum_{e \in T_\beta} P_i(e) \times p_\beta(e)$ for $i = 1, \dots, n$.

In the spirit of conserving letters, the original letter P that denotes payoffs assigned to the terminal paths is recycled here.

Finally, pure strategies are defined as a special case of behavioral strategies.

Pure strategies: β_i , such that $\beta_i(I_i^k)$ is always degenerate, i.e., picks only one action with certainty, is called a *pure strategy* and is denoted by π_i . For pure strategies, the notation π is used in place of β and Π is used in place of B .

Decomposition of strategies: The definitions offered below introduce certain partial strategies or strategy profiles for G and β :

- β^H is β reduced to H if $\beta^H = \cup_{i=1}^n \beta_i^H$;
- β_i^{-H} is a complement of β_i with respect to H if $\beta_i^{-H} = \beta_i - \beta_i^H$;
- $\beta^{-H} = \cup_{i=1}^n \beta_i^{-H}$ is a complement of β with respect to H ;
- B_i^{-H} is the set of all β_i^{-H} for all $\beta_i \in B_i$;
- $B^{-H} = \times_{i=1}^n B_i^{-H}$.

Let $\delta_i^H : B_i \rightarrow B_i^H \times B_i^{-H}$ denote the *decomposition function for player i* , which assigns to β_i its reduced strategy β_i^H and its complement β_i^{-H} . The *decomposition function* $\delta^H : B \rightarrow B^H \times B^{-H}$ is defined as $(\delta_i^H)_{i=1}^n$. The following simple but useful result holds for every game G and its subgame H :

Lemma 2. (a) For every i, δ_i^H is 1-1 and onto;

- (b) δ^H is 1-1 and onto.

Lemma 2 allows us to define the function of the *concatenation* of strategies that is the inverse of decomposition: For every subgame H of G and every pair of partial strategy profiles $\beta^H \in B^H$ and $\beta^{-H} \in B^{-H}, \sigma^H(\beta^H, \beta^{-H}) = \beta^H \cup \beta^{-H}$ is such that $\sigma^H \delta^H(\beta^H) = \beta^H$. Moreover, $\sigma^H = (\sigma_i^H)_{i=1}^n$, where every σ_i^H is the inverse of a respective δ_i^H . Similar to δ^H , both σ^H and all its all components σ_i^H are 1-1 and onto.

The final two definitions of this section introduce two familiar equilibrium concepts [8,31]. For any game G and the strategy profile $\beta \in B$, the equilibrium conditions for β are stated as follows:

Nash equilibrium (NE): For every $i \in N, \beta_i \in \text{ArgMax}_{t_i \in B_i} P_i(\beta_{-i}, t_i)$;

Subgame perfect equilibrium (SPE): For every subgame H, β^H is an NE in H .

Analogous definitions hold when all considered strategies are pure.

3. Decomposition of Strategies

The notations s_i, \mathbf{s}, S_i, S , etc. are used to denote the strategies, strategy profiles, strategy spaces, joint strategy spaces, etc. that are either pure or behavioral in order to process both cases simultaneously.

Lemma 2 guarantees that the operations δ and σ are well defined and that they bring unique outcomes within the same family of strategies. Obviously, the family of pure strategies has the same property. Moreover, the definition of finite support of every strategy guarantees that the outcomes of δ and σ have finite support.

The profiles \mathbf{s} or \mathbf{s}^G denote any strategy profiles in G , and \mathbf{s}^H and \mathbf{s}^{-H} (or \mathbf{s}^{G-H}) denote their decomposition with respect to its subgame H .

The sets $T_{s(G-H)}$ and $T_{s(H)}$ denote all terminal paths from T_s that do not include the root of H , ϕ , or that do include ϕ , respectively:

$$T_{s(G-H)} = \{e \in T_s \text{ such that } \phi \notin e\};$$

$$T_{s(H)} = \{e \in T_s \text{ such that } \phi \in e\}.$$

Lemma 3 states that the payoff in any game G from any strategy profile \mathbf{s} is the sum of the payoffs from all terminal paths that do not include ϕ and the payoff of \mathbf{s} reduced to H multiplied by the probability of reaching ϕ .

Lemma 3. $P^G(\mathbf{s}) = p_{sG}(\phi)P^H(\mathbf{s}^H) + \sum_{e \in T_{s(G-H)}} p_{sG}(e)P^G(e).$

Upgame: For any game $G = \langle \Psi, N^0, \{T_i\}_{i \in N^0}, I, A, h, P \rangle$, an *upgame* of G (with respect to a subgame H) is any game $F = \langle \Psi', N^0, \{T'_i\}_{i \in N^0}, I', A', h', P' \rangle$ if (a) Ψ' is a subtree of Ψ such that ϕ , the root of H in G , and all nodes that follow ϕ are substituted with a terminal node ϕ in F and a payoff vector $P^F(\phi)$ that is of the same dimension as the payoffs in G and (b) the players are unchanged and $\{T'_i\}_{i \in N^0}, I', A', h'$, and P' are the restrictions of f, I, A, h , and P to Ψ' , respectively (with ϕ excluded from the restriction).

The demonstration that such restrictions define a game is straightforward. In a similar fashion, we can substitute any non-empty set of disjoint subgames of G . Every game resulting from such an operation is also called an upgame.

(s, Θ)-upgame: F is an upgame of G with respect to a strategy profile \mathbf{s} and a non-empty set of the disjoint subgames $\{H_\theta\}_{\theta \in \Theta}$ of G , where ϕ_θ is a root of H_θ if $P^F(\phi_\theta) = P^{H_\theta}(\mathbf{s}^{H_\theta})$. If for each H_θ , \mathbf{s}^{H_θ} is SPE, then F is called a *perfect upgame*. If $|\Theta| = 1$, the notation is *(s, H)-upgame*.

An upgame is obtained when we substitute the roots of disjoint subgames from a set Θ with arbitrary payoff vectors. When such vectors result from a strategy profile \mathbf{s} acting in the respective subgames, it is an *(s, Θ)-upgame*. It becomes a perfect upgame when every \mathbf{s}^{H_θ} is SPE in H . Note that the classical backward induction prunes game trees by building perfect upgames one at a time.

It is useful to note a few facts. For a family of subgames $\{H_\theta\}_{\theta \in \Theta}$, no perfect upgames may exist or, alternatively, there may be multiple perfect upgames. If a subgame H_θ has no SPE, then no perfect upgame exists for H_θ and no SPE exists for the entire game.

Let F be an *(s, H)-upgame* of G for any strategy profile $\mathbf{s} \in S$. The following Lemma states a simple relationship between the payoffs in G and F .

Lemma 4. $P^G(\mathbf{s}) = P^F(\mathbf{s}^{-H}).$

The next result characterizes the fundamental aspect of pruning a game. Since the concatenation and reduction of strategies will be applied to the subgames of subgames, we need additional notations:

\mathbf{s}^{HJ} is a strategy profile \mathbf{s} reduced to a subgame H and then further reduced to J —a subgame of H ; additionally, \mathbf{s}^{H-J} is a complement of \mathbf{s}^J in H . Similar notations are applied to individual strategies and payoff profiles.

For any game G and any of its proper subgames H and for any strategy profile \mathbf{s} , let F be the *(s, H)-upgame* of G .

Theorem 1. (*decomposition*): *The following conditions are equivalent:*

- (a) \mathbf{s} is an SPE for G ;
 (b) \mathbf{s}^H is an SPE for H , and \mathbf{s}^{-H} is an SPE for F .

The Decomposition Theorem states that every SPE can be obtained by the concatenation of two SPE subgame-upgame profiles and that every concatenation of two SPE profiles produces an SPE.

Agenda: Consider the graph that includes the roots of all subgames of G , that has the same root as G , and of which the successor relationship is imported from Ψ . It is clear that such a graph is a game tree. By its obvious association with voting models, it is called the *agenda* of G , and the set of all agenda nodes is denoted by T_A .

Subgame level: For a subgame H of game G with a root ϕ , the *level* of H is the total number of nodes that are followed by ϕ in the agenda of G (including both τ and ϕ).

It is clear that the level of any subgame here is a positive integer.

Lemma 5. For any game G , any positive integer k , and any two different subgames H and J of G of level k , the sets of nodes of H and J are disjoint.

Lemma 5 implies that we can substitute any set of subgames of the same level with payoffs of the appropriate dimension and obtain an upgame of F . Removing all subgames at the same level is convenient, and this assumption appears in many applications. However, it is sufficient to assume that all removed subgames are disjoint.

Theorem 1 is now extended to any set of disjoint subgames.

For any game G , any subset of disjoint subgames $\{H_\theta\}_{\theta \in \Theta}$ of G , and any profile \mathbf{s} in G , let F be the (\mathbf{s}, Θ) -upgame of G .

Theorem 2. (*simultaneous decomposition*): The following conditions are equivalent:

- (a) \mathbf{s} is an SPE for G ;
 (b) \mathbf{s}^F is an SPE for F , and for all $\theta \in \Theta$, \mathbf{s}^{H_θ} is an SPE for H_θ .

4. Generalized backward Induction (GBI) Algorithm

Theorem 2 legitimizes a general procedure of backward induction for any game and the pure or behavioral strategies of finite support and finite crossing. GBI proceeds up the game tree by concatenating partial SPE strategy profiles in consecutive disjoint sets of subgames.

Let us fix the game G .

Pruning sequence: The sequence of pruning $\{\Theta_j\}_{j=1}^l$ is a partition of T_A , the set of agenda nodes, where l is a positive integer, $l \geq 2$, such that for all $k, m \in \{1, \dots, l\}$, if $\chi \in \Theta_k$ follows $\psi \in \Theta_m$, then $k \leq m$.

The pruning sequence denotes the order of removing the subgames, with Θ_j denoting the roots of the subgames removed in step j . The condition imposed on $\{\Theta_j\}_{j=1}^l$ asserts that a subgame J of a subgame H is pruned before or, simultaneously with, H .

Let us consider the agenda and all possible pruning sequences for a simple example of pure coordination with perfect information (see Figure 2).

The agenda of pure coordination includes three nodes: A1, B1, and B2. There are six possible pruning sequences: {B1}-{B2}-{A1}; {B2}-{B1}-{A1}; {B1, B2}-{A1}; {B1}-{B2, A1}; {B2}-{B1, A1}; and {B1, B2, A1}. According to the first two sequences, single subgames are pruned; in the remaining sequences, pruning includes removing two or three subgames at the same time. The condition imposed on the pruning sequence guarantees that A1 is pruned in the last step, possibly with other nodes.




Figure 2. Pure coordination with perfect information.

Finding a backward induction solution (BIS) begins with the entire game $G = G_1$. In every step of pruning, a new perfect upgame G_{j+1} of G_j is created according to the pruning sequence $\{\Theta_j\}_{j=1}^l$. A BIS exists if, for at least one sequence of pruning, such a sequence of perfect upgames can be found:

Backward induction solution (BIS): Strategy profile s is a BIS according to the pruning sequence $\{\Theta_j\}_{j=1}^l$ if (a) there is a set of games $\{G_j\}_{j=1}^l$ such that $G_1=G$ and for $j = 1, \dots, l - 1$, G_{j+1} is a perfect (s, Θ) -upgame of G_j .

In other words, s is BIS if we can prune a game using s in such a way that, at every stage, s is an SPE in the removed subgames and s is also an SPE in the the final upgame that results from pruning.

Let us go back to pure coordination. Consider the pruning sequence B1-B2-A1. After the removal of the first subgame, we replace the root of this subgame B1 with the SPE payoff (1,1). The set of partial SPEs includes one partial strategy a : $S_1^{SPE} = \{(a)\}$. In the second step, after removal of B2, the new partial strategy d is concatenated with the previously obtained partial strategy and $S_2^{SPE} = \{(ad)\}$. In the final step, both L and R are the SPE in the final perfect upgame and they can be concatenated with the previously obtained partial strategy: $S_3^{SPE} = \{(ad; L), (ad; R)\}$.

We have now the tools that allow us to examine the relationship between S^{BIS} and S^{SPE} . For a fixed game G and a set of strategy profiles (behavioral or pure) S , let us denote the subset of all BISs with S^{BIS} and the subset of all SPEs by S^{SPE} . Using our definition of BIS as resulting from *any* sequence of pruning, if s is SPE, then for $l = 2$, by Theorem 2, s is also BIS. Conversely, if s is a BIS, then we can find a pruning sequence $\{\Theta_j\}_{j=1}^l$ that satisfies the conditions from the definition of BIS. Theorem 2 applied $l - 1$ times guarantees that s is SPE. The relationship between subgame perfection and backward induction can now be stated formally. It is straightforward:

Corollary 1. For any game G , $B^{SPE} = B^{BIS}$ and $\Pi^{SPE} = \Pi^{BIS}$.

A simple consequence of Corollary 1 (in combination with Theorem 2) is that if s is BIS with one pruning sequence, then it must be BIS with any pruning sequence. The only differentiating factor is the convenience of using one sequence over another.

The following algorithm describes finding all SPEs:

Generalized Backward Induction (GBI) Algorithm:

1. *Initial pruning:* Set a pruning sequence $\{\Theta_j\}_{j=1}^l$ with the subgames $\{H_\theta\}_{\theta \in \Theta_j}$ of consecutive upgames of G pruned in step j . Set the initial set of partial strategy profiles S_1^{SPE} , defined as the set of all partial strategy profiles in G that are SPE for all subgames of G with roots from Θ_1 , i.e., for $\{H_\theta\}_{\theta \in \Theta_1}$.
2. *Verification of partial SPEs:* In step j , where $1 \leq j \leq l$, the procedure generated S_j^{SPE} .

- If $j = l$, set $S^{SPE} = S_l^{SPE}$ and stop. Otherwise:
- If $j < l$ and $S_j^{SPE} = \emptyset$, then set $S^{SPE} = \emptyset$ and stop.
- If $j < l$ and $S_j^{SPE} \neq \emptyset$, then go to 3.

At step j , set S_j^{SPE} is the set of all partial strategy profiles that were obtained for the pruned subgames up to the level Θ_j . If S_j^{SPE} is empty for any j , this implies that set S^{SPE} is also empty. A non-empty S_j^{SPE} may include more than one partial strategy profile.

3. *Concatenation*: Let us denote the elements of S_j^{SPE} by \mathbf{s}_j^k , for $k \in K$. For every $\mathbf{s}_j^k \in S_j^{SPE}$, perform the following procedure for every $\mathbf{s}^{\Theta_{j+1}}$, a partial strategy profile for all subgames $\{H_\theta\}_{\theta \in \Theta_{j+1}}$. Exactly one of the following must hold:

- (i) $\mathbf{s}^{\Theta_{j+1}}$ is an SPE for all $\{H_\theta\}_{\theta \in \Theta_{j+1}}$. In such a case, include $\mathbf{s}_j^k \cup \mathbf{s}^{\Theta_{j+1}}$ in S_{j+1}^{SPE} ;
- (ii) $\mathbf{s}^{\Theta_{j+1}}$ is not an SPE for at least one of $\{H_\theta\}_{\theta \in \Theta_{j+1}}$. In such a case, discard $\mathbf{s}_j^k \cup \mathbf{s}^{\Theta_{j+1}}$.

In every step of concatenation, each partial strategy profile from S_j^{SPE} is checked against each partial strategy profile for the next set of subgames. If the concatenation of both profiles produces an SPE, it is included in the next set of partial SPEs, S_{j+1}^{SPE} . Otherwise, it is discarded.

- 4. Increase j by one and go back to 2.

5. Examples

Below, I discuss five examples of applying the GBI algorithm. At every step, instead of moving the payoff vectors up—which is the ordinary BI procedure—the set of partial SPEs is created. The final set includes all SPEs. Also, note that in all five examples, unlike ordinary BI, GBI prunes many subgames simultaneously.

Although GBI helps to solve some games, one should mention its restrictions. A rather obvious limitation is that if a game’s agenda is a singleton—meaning that it has no proper subgames—then GBI offers no benefit, as no pruning is possible. Moreover, for an infinitely repeated game, such as the Prisoner’s Dilemma [32,33], replacing an infinitely long subgame or subgames with an SPE is essentially equivalent to figuring out an SPE for the entire game. The usefulness of the method depends mostly on the structure of the agenda.

5.1. Complex Finite Games with Perfect Information

In this classic case, the agenda is identical to the game tree minus the endnodes. Finding an SPE in every subgame is equivalent to finding the best move (or the best moves) of a player. If payoffs at some stage are identical, one may obtain many SPEs. The existence of a BIS for finite games was proved by Kuhn’s Corollary 1 [6] (p. 61).

Even a finite game—if it is complex—may benefit from the GBI algorithm.

Example 1. *Pick 100*: Alice starts the game by picking any positive integer $x \in \{1, 2, \dots, 10\}$. Next, Bob picks a greater integer y , such that $(y - x) \in \{1, 2, \dots, 10\}$. Then, Alice picks z , such that $(z - y) \in \{1, 2, \dots, 10\}$, etc. The game ends when someone picks at least 100. The winner’s payoff is 1; the loser’s payoff is -1 .

The game is finite, but close pruning required by classic BI is unrealistic due to the enormity of the game tree. For instance, there are 512 subgames with their roots labeled with 10. This is the total number of different paths that lead from 0 to a subgame that begins with 10. With the label increasing, the number of corresponding subgames increases quickly.

What is the “solution” to this game? It can be described intuitively (the solution is presented at the end of the example), but its relation to SPE is unclear. Moreover, the calculation of all SPEs is complicated. The GBI helps by applying simultaneous pruning of large sets of disjoint subgames.

The sequence of pruning includes all maximal subgames labeled with certain numbers, as described below.

First step: Starting from 90–99, a player’s best action is to pick 100. Thus, we have to replace all maximal subgames with roots of at least 90 with the corresponding payoffs (1, −1) or (−1, 1), depending on whether the player is Alice or Bob. There are 10 types of subgames labeled 90–99 per player, but the number of subgames of the same type is very large since they can be reached via many different paths. Figure 3 shows the three simplest types of subgames pruned in Step 1 that correspond to the previously picked numbers of 99, 98, and 97 (players and player payoffs are not represented). SPE strategy profiles are marked in bold. Set S_1^{SPE} includes simply all strategy profiles, “pick 100 in all subgames labeled 90–99, and their subgames.” Exactly one partial strategy profile satisfies this requirement.




Figure 3. The tree types of subgames pruned in Step 1.

Second step: Since all subgames starting with 90–99 were pruned, the greatest remaining number is 89. The player who picks 89 wins, because the other player is now forced to pick a number between 90 and 99. Thus, in this step, we prune all subgames that begin with 89. The losing player, who must choose a number between 90 and 99, has the only available payoff of −1. Thus, any partial strategy profile is SPE in all subgames. When such profiles are concatenated with the single profile from S_1^{SPE} , we obtain the following second set of partial SPEs that can be defined informally as follows:

$$S_2^{SPE} = \{s : \text{from 89, pick anything; from 90–99 and their subgames, always pick 100}\}$$

Third step: Now, picking 89 means winning the game since, in the new upgame, all endnodes labeled with 89 are terminal and offer the payoff of 1. We can prune all maximal subgames that start with a number between 79 and 88. The single best action for both players is to pick 89. The SPE partial strategy profile satisfies the condition “always pick 89”.

Next steps: Similarly, the player who picks 79–88 loses; the player who picks 78 wins, since the next player must pick 79–88 and so on. In the last step, it turns out that Alice can pick 1, the first number in the winning sequence. Her winning sequence of moves is concatenated from 20 levels of pruning with the picks of 1, 12, 23, 34, 45, 56, 67, 78, 89, and 100 regardless of Bob’s choices. An interesting property of the game Pick 100 is that, despite the game’s complexity, defining Alice’s ten crucial types of actions guarantees her a winning path. There is a large number of SPEs:

$S^{SPE} = \{s : \text{both players pick one of the numbers 1, 12, 23, 34, 45, 56, 67, 78, 89, and 100 if they can do that and anything if they cannot}\}$

Certain strategies that are not part of any SPE can guarantee Alice winning as long as she chooses the sequence of ten magical numbers. For instance, since she is not starting the game with 2, then whatever she would choose in the subgame following 2 would not upset her winning path that starts with 1.

Now, we can define a “solution” to Pick 100 and similar games as a much simpler object than an SPE: A solution is any minimal set of actions that guarantees some player a winning path.

5.2. Continuum of Actions and Perfect Information

Examples of nonfinite games of perfect information include various games of fair division [34,35], the Romer–Rosenthal Agenda Setter model [36], and the Ultimatum game. The algorithm for such games closely resembles classic backward induction. Below, I will demonstrate how GBI solves the Agenda Setter model.

Example 2. Romer–Rosenthal Agenda Setter model.

Two players, Agenda Setter A and Legislator L , have Euclidean preferences in the issue space $[0, 3]$ and the ideal points 0 and 2, respectively. The status quo is $q = 3$. The dynamics are as follows:

Stage 1: A proposes a policy $x \in [0, 3)$.

Stage 2: L chooses the law from $\{x, q\}$.

The Agenda Setter model defines a unique game G with the following components:

$$\Pi_A = [0, 3); \Pi_L = \{X : X \subset [0, 3)\}$$

Every strategy profile includes (x, X) , where x is a policy proposed by A , and set X represents all policies that L would accept.

$$P_A(x, X) = \begin{cases} -x & \text{if } x \in X \\ -3 & \text{if } x \notin X \end{cases}$$

$$P_L(x, X) = \begin{cases} -|x - 2| & \text{if } x \in X \\ -1 & \text{if } x \notin X \end{cases}$$

The payoff of A is the negative distance between A 's ideal point, 0, and the new law. It is equal to $-x$ if x is accepted by L , and it is equal to -3 if L rejects x . L 's payoff is defined similarly.

Both GBI steps correspond to one subgame level. The adopted pruning sequence simultaneously removes all subgames at the same level. Here, we are only interested in pure strategies.

Step 1: At level two, there is a continuum of subgames parametrized by the issue space $[0, 3)$. When $x \in [0, 3)$ is proposed, L has two options: to accept it or to reject it (which implies that q is accepted). For a subgame with its root at x , the best actions for L are the following:

If $x < 1$, reject x ;

If $x > 1$, accept x ;

If $x = 1$, reject or accept x .

Applying simultaneous pruning to level 2 brings our first set of partial SPEs, with two partial SPE profiles equal to two strategies of L :

$S_1^{SPE} = \{[1, 3), (1, 3)\}$, i.e., “accept every offer not smaller than 1” and “accept every offer greater than 1”.

The two partial strategy profiles produce two perfect upgames G_1 and G_2 with player A and their strategy space $[0, 3)$, where their payoffs differ only for $x = 1$ and are defined as follows:

$$G_1 : P_A(x) = \begin{cases} -3 & \text{if } x < 1 \\ -x & \text{if } x \geq 1 \end{cases}$$

$$G_2 : P_A(x) = \begin{cases} -3 & \text{if } x \leq 1 \\ -x & \text{if } x > 1 \end{cases}$$

Step 2: Now, we consider both partial strategy profiles from S_1^{SPE} , i.e., $[1, 3)$ and $(1, 3)$.

Partial strategy profile $[1, 3]$: The unique best action for A in G_1 is $x = 1$ since it maximizes A 's payoff among the options that are acceptable to L with $P_A(1; [1, 3]) = -1$. When this move is concatenated with $[1, 3]$, the resulting SPE for the entire game is $(1; [1, 3])$;

Partial strategy profile $(1, 3)$: There is no best action for A in G_2 since L 's set of acceptable payoffs $[-3, -1]$ does not include its upper bound. Partial profile $(1, 3)$ is discarded.

Solution: There is a unique SPE in G , $S^{SPE} = \{(1; [1, 3])\}$. In the SPE, A offers 1 and L accepts.

5.3. Finitely Repeated Games

An interesting case arises when a finite game is finitely repeated.

First, let us assume that G has precisely one equilibrium in behavioral strategies. Let G^k be G repeated k times, $k \geq 2$. When close pruning is applied, there is precisely one possible SPE in every removed subgame of G^k that corresponds to the SPE in G . When all subgames of the same level are pruned, the resulting game is G^{k-1} plus a fixed payoff adjustment for all players equal to the equilibrium payoff in G , which does not affect the equilibrium. This fact implies the following result (s denotes either the pure or behavioral strategy):

Corollary 2. For any finite game G that has a unique SPE s and for any integer $k \geq 2$, G^k has exactly one SPE that is equal to the repeated concatenation of s .

A consequence of Corollary 2 is the well-known fact that such finitely repeated games as the Prisoner's Dilemma or Matching Pennies have exactly one SPE in pure and behavioral strategies, respectively.

When a one-shot game has many SPEs, different equilibria may contribute different payoff vectors to the upgames. Nevertheless, GBI may simplify the calculations.

Example 3. Twice-repeated pure coordination.

In one-shot pure coordination (PC), Alice and Bob simultaneously choose one of their two strategies (refer to the game shown in Figure 2 plus imperfect information). There are two levels in the twice-repeated PC. Our sequence of pruning once again coincides with the levels.

Step 1: There are four subgames at level two, with three NEs in each subgame: $(L; L)$, $(R; R)$, and the NE in completely mixed strategies $(\frac{1}{2}, \frac{1}{2})$. This produces 81 partial strategy profiles that are NEs; 16 of them are in pure strategies. The sets of partial SPEs in pure and behavioral strategies are defined as follows, respectively:

$$\Pi_1^{SPE} = \{y, z, v, w \text{ such that } y, z, v, w \in \{(L; L), (R; R)\}\};$$

$$B_1^{SPE} = \{y, z, v, w \text{ such that } y, z, v, w \in \{(L; L), (R; R), (\frac{1}{2}; \frac{1}{2})\}\}.$$

Step 2: At this point, it makes sense to separate the cases of pure and behavioral strategies.

Pure strategies: In this easy case, every partial profile from Π_1^{SPE} obtained in Step 1 adds exactly 1 to the payoff in the upgame. Thus, for every partial strategy profile obtained in Step 1, there are two SPEs in the upgame, i.e., the coordination on L or on R . Consequently, there are 32 SPEs in pure strategies that can be described as follows:

$$\Pi^{SPE} = \{x, y, z, v, w \text{ such that } x, y, z, v, w \in \{(L; L), (R; R)\}\}.$$

Behavioral strategies: Each perfect upgame that results from pruning all four subgames has exactly three NEs. In every upgame, players receive the payoffs of $1\frac{1}{2}$ or 2 for coordinating their strategies and 1 or $\frac{1}{2}$ for discoordination. Thus, every perfect upgame is either PC or a variant of asymmetric coordination having exactly two NEs in pure strategies and one NE in completely mixed strategies. Each of the 81 partial strategy profiles from Step 1 can be concatenated with the three NEs in the upgame. The total number of SPEs in behavioral strategies is 243.

By not enumerating all behavioral SPEs, I will conserve space for the next example.

5.4. Parametrized Games in Political Economy

In the field of political economy, typical models are often represented as parametrized families of games (in short, parametrized games). SPE is often an appropriate solution concept. As the next example, I apply a simplified Auto-Lustration (AL) model [37], which is a more complex version of the Agenda Setter model, motivated by the surprising behavior of postcommunist parties in the 1990s. When such parties returned to power in some Central European countries, just before the end of their terms, they started legislating light auto-lustration. In other words, they were punishing their own members for being former supporters of communism! The term “lustration” signifies some punishment imposed on former functionaries or on the secret informers of communist regimes such as making their names public or blacklisting them from certain public offices.

Example 4. Auto-Lustration.

A postcommunist party (P) and an anti-communist party (A) have Euclidean preferences over the lustration space $[0, 1]$. P 's ideal amount of lustration is zero; for A , the ideal amount is one. There is also a smaller moderate party, M , whose ideal point is slightly tilted to the right $m > \frac{1}{2}$. M has no chance of winning a majority and will join a postelection coalition only with A .

The game unfolds as follows (recall that the dynamics are simplified and that the choices are restricted):

Period 1: P is the ruling party and can choose between the status quo of no lustration (0) and a moderate amount of lustration, m , which is the ideal point of M ;

Period 2: Parliamentary elections take place and a new parliamentary median is elected with the following probabilities (for the extreme parties A and P , being a median is the equivalent of winning the majority):

A : p_1 ;

M : p_2 ;

P : $(1 - p_1 - p_2)$.

Period 3: If P is the new median, then they do not introduce any new legislation, since any change, including reverting to 0, would result in an unacceptable loss of credibility in the eyes of their electorate.

If A is the new median, then they can choose any law.

If M is the new median, they coalesce with the larger partner A . M 's approval is necessary for any new law. If the existing legislation is 0, A proposes new legislation and then M must either accept or reject it. If the existing legislation is m , A proposes no new legislation, anticipating that it will not receive M 's support.

Period 4 (only if M is the new median and the existing legislation is 0): M either accepts A 's proposal or the status quo prevails.

Figure 4 shows the AL game with the issue space or the outcomes in place of the payoffs.

The game is parametrized by two probabilities, p_1 and p_2 , and the position of the moderate party, m .

The player strategy spaces are as follows:

$$\Pi_P = \{0, m\};$$

$$\Pi_A = [0, 1]^3;$$

$$\Pi_M = 2^{(0,1]}.$$




Figure 4. Simplified structure of the Auto-Lustration game.

P's strategy includes the initial choice between light lustration m and no lustration. *A*'s strategy involves three scenarios of legislation, depending on what *P* decided earlier and whether the new median is *A* or their coalition partner *M*. When *M* is the new median and *P* did not change the status quo, *M* must decide whether to accept *A*'s proposal or to keep 0. Thus, *M*'s strategy is, like in the Agenda Setter model, a subset of acceptable policies that are preferred to the status quo.

The payoffs are the negative distances of the final lustration law from the parties' ideal points. We assume that the parties are risk neutral. The game is fairly complex and asymmetric, but the calculation of SPE is straightforward with GBI. The subscripts in the partial strategy profiles denote the player who plays that particular partial strategy.

Step 1: In its terminal set of subgames, due to its position slightly tilted to the right, *M* prefers anything to the status quo 0, i.e., $S_1^{SPE} = \{(0, 1]_M\}$.

Step 2: The SPEs in the three subgames of Period 3 are as follows:

- 1 and 2: If *A* wins the majority, they propose the harshest lustration 1.
- 3: If *M* is the median and the status quo is 0, *A* proposes 1, since they know that *M* prefers 1 to 0.

Applying the above order of listing the subgames, $S_2^{SPE} = \{(1_A 1_A 1_A; (0, 1]_M)\}$ (i.e., the set of partial SPEs includes exactly one partial strategy).

Step 3: *P* chooses between the payoffs in the perfect upgame, reducing the game to a choice between 0 and m . Introducing a light lustration law is strictly preferable if the expected payoff from playing m is higher than it is from playing 0:

$$-p_1 - p_2 < -p_1 - p_2 m - (1 - p_1 - p_2)m \tag{1}$$

After simplification, $p_1 m + p_2 - m > 0$

This is the equilibrium condition for m to be concatenated to the previously obtained partial strategy profile and to form a unique SPE: $(1_A 1_A 1_A; m_P; (0, 1]_M)$. When the inequality is reversed, 0 is the newly partial strategy profile and the unique SPE is $(1_A 1_A 1_A; 0_P; (0, 1]_M)$; with equality, both 0 and m can be concatenated to form two SPEs.

In Hungary and Poland, *P* won the majority in the 1994 and 1993 elections only because the rightist parties were fragmented in those early elections and were unable to form a unified bloc. It was practically certain that, in the new elections, either *A* or *M* would win. Note that with $p_1 + p_2 = 1$ and $p_2 > 0$, the inequality is satisfied. Light auto-lustration was a sensible strategy as insurance against the harsher punishment by the new government. In fact, in both Hungary and Poland, *A* won the next

election but M became the new median, and despite many attempts, A was unable to strengthen the existing lustration law significantly [38].

5.5. Weakly Undominated Strategies in Subgames

A typical voting game may involve a large number of SPEs that are unreasonable. The GBI algorithm can be modified in order to eliminate such unreasonable equilibria. While the modification explained below is ad hoc, it is worthwhile to examine how it works in a specific game.

Example 5. Roman Punishment.

Farquharson [39] and Riker [40] analyzed the apparently first recorded case of a voting manipulation. The letter of the manipulator, Pliny the Younger, reported the story of a decision made by the Roman Senate. Three groups of senators were deciding the fate of a freedman, who was possibly involved in the death of a Roman consul. According to the Roman judiciary agenda, which was clearly quite different from the modern court procedures, the senators needed to decide first whether the freedman deserved death or not. A negative answer would trigger the next decision, whether to banish or acquit the freedman. The game analyzed below should have taken place in the Roman Senate according to its normal agenda. However, it did not happen since Pliny persuaded the senators to use a simpler plurality agenda in which they voted over the three options simultaneously.

Voting rule: simple majority, no abstention

Round 1 alternatives: d (death) or n (no death)

Round 2 alternatives (if n wins in the first round): b (banish) or a (acquit)

Players (the names correspond to the player top alternatives): A (acquitters), B (banishers), and D (death penalty supporters)

Player preferences:

A : a preferred to b preferred to d

B : b preferred to a preferred to d

D : d preferred to b preferred to a

Let us convert the player preferences into payoffs in the following way:

3—the top alternative for every player

2—the second-best alternative for every player

1—the worst alternative for every player

Figure 5 depicts the extensive game representing voting according to the Roman agenda.

The game looks simple but the number of SPEs is staggering. H_4 depicted in Figure 5 has three SPEs, i.e., $S_{H_4}^{SPE} = \{(b; b; b), (a; a; a), (a; b; b)\}$. When all players vote identically, either $(b; b; b)$ or $(a; a; a)$, their vote is an NE (and also SPE) since no deviation of one player can change the outcome of voting. The third SPE, $(a; b; b)$, is when all voters vote for their preferred alternative, i.e., when they choose their strategies that are weakly dominant in H_4 .

The remaining three subgames H_1 , H_2 , and H_3 have a structure that is identical to H_4 and have three SPEs each. Thus, when we prune all four subgames, the set of partial SPEs, S_1^{SPE} , includes $3^4 = 81$ partial strategy profiles. There are $2^4 = 16$ possible upgames since in each subgame, either a or b , can be the equilibrium outcome with the corresponding payoff vectors $(3, 2, 1)$ or $(2, 3, 2)$. Thus, every subgame may be replaced either with $(3, 2, 1)$ or $(2, 3, 2)$.

At this moment, I abandon calculating all SPEs since the number is large and all of them except for one have a fatal flaw. Namely, except for the partial strategy profile $(aaaa; bbbb; bbbb)$, in all other strategy profiles, at least one player plays a weakly dominated strategy in at least one subgame, i.e., votes for the second-best alternative. Let us consider what happens in the perfect upgame resulting from this special profile that steers clear of weak domination $(aaaa; bbbb; bbbb)$ (see Figure 6).

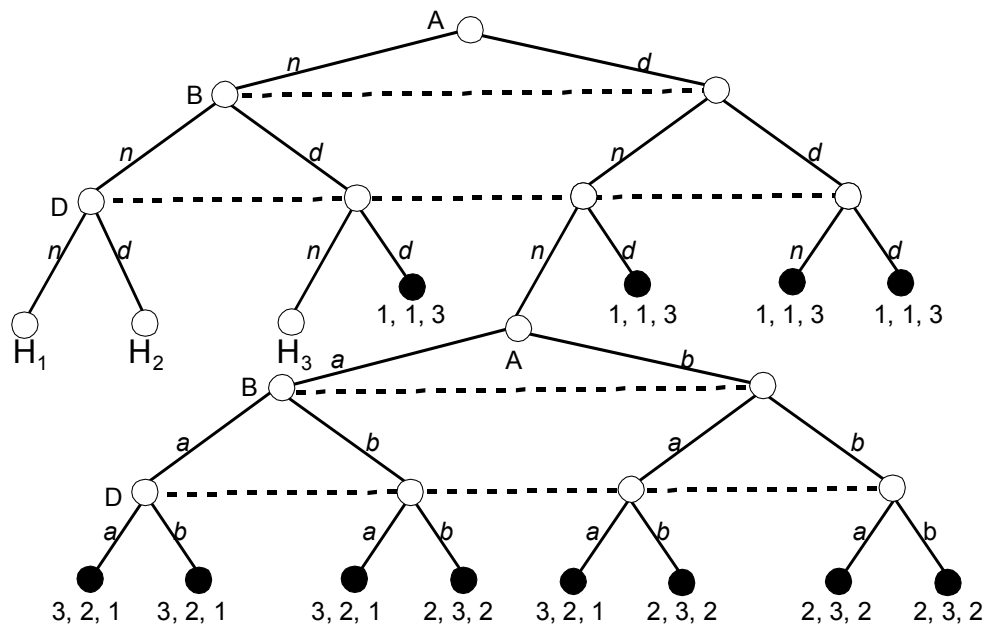


Figure 5. Roman Punishment: Out of four subgames, only H_4 is shown.




Figure 6. The perfect upgame of Roman Punishment corresponding to $(aaaa, bbbb, bbbb)$.

A quick calculation shows that the players in our upgame have three SPEs similarly to the four removed subgames. In addition to $(n; n; n)$ and $(d; d; d)$ that involve at least one weakly dominated strategy, there is an SPE in weakly dominant strategies $(n; n; d)$. The concatenation of this SPE with the previously obtained $(aaaa; bbbb; bbbb)$ brings an SPE $(naaaa; nbbbb; dbbbb)$ that has an additional property that in all removed subgames and in the final upgame, no player plays a weakly dominated strategy.

It is straightforward to show that, in Roman Punishment, the strategy profile $(naaaa; nbbbb; dbbbb)$ is a unique perfect equilibrium (PE). All partial strategy profiles are also unique PEs in the respective subgames. Interesting questions arise: (1) What are the general conditions for the concatenation of partial PEs to produce a global PE? (2) Can we obtain all PEs that way? and (3) Does the final result depend on the pruning sequence? Moreover, since any equilibrium concept refining the SPE can be applied at all stages of GBI, similar questions arise for all other refinements.

6. Conclusions

The contributions of this paper include the following: (a) the axiomatization of infinite games; (b) a demonstration that BIS and SPE are equivalent for such games for pure and behavioral strategies of finite support and crossing; and (c) the provision of an algorithm for solving certain games.

Infinite games that I consider may have imperfect information, infinite action sets, and an infinite horizon. Informally, the algorithm operates as follows (the formal presentation is found in Section 4):

1. Identify the game's agenda, i.e., the tree consisting of all roots of the game's subgames that is ordered by the relation of the successor imported from the game tree. Set the pruning sequence.
2. Prune any subgame according to the pruning sequence and substitute its root with the subgame's SPE payoffs. The procedure of substitution may be conducted simultaneously for any subset of agenda nodes as long as the corresponding subgames are pairwise disjoint.
3. Concatenate all partial strategy profiles resulting from the substitution.
4. If at any point one receives an empty set as SPE for the subgame, this would mean that there is no SPE compatible with the set of previously selected SPEs for the subgames.
5. In order to find all SPEs, one needs to try all possible substitutions of the subgames with SPE payoffs.
6. One stops at the root of the game.

The examples discussed above illustrate the application of the algorithm to complex games that involve simultaneous pruning of large numbers of subgames (Example 1), continuum of actions (Example 2), behavioral strategies and imperfect information (Example 3), and parametrized games (Example 4). Example 5 shows an ad hoc modification of GBI that allows to find the unique perfect equilibrium in a game that has a large number of unreasonable SPEs.

Three open problems deserve a further comment:

1. *Extending the results:* An obvious open question is whether the results for behavioral strategies of finite support and crossing can be generalized to all "rough" behavioral strategies. Attacking this question would demand leaving the comfortable world of finite probability distributions and using measure theory in the spirit of Aumann's [30] pioneering contribution. The framework presented in this paper goes around measure-theoretic difficulties by assuming finite support and crossing. Both assumptions imply that the total number of terminal paths that count for calculating payoffs is finite for every strategy profile. It is easy to identify the places in the proofs where this fact is used. A natural question, then, is whether the equivalence can be extended.
2. *Axiomatization of noncooperative game theory:* The general axiomatic framework applied in the present paper encompasses more games than the classical approaches of von Neumann [4] and Kuhn [6]. When game theory was born, it seemed natural to consider only finite games; nonfinite games rarely appeared in the literature. Today, we routinely go beyond the limitations of finite games, either with a continuum of strategies that represent quantity, price, or position in the issue space or with the infinite repetition of a game. I believe that contemporary game theory deserves sound axiomatic foundations that can cover infinite games. This would lead toward a more unified and complete discipline. Concepts that were axiomatically analyzed for finite games, such as Kreps and Wilson's (1982) sequential equilibrium, seem to be obvious targets for a more comprehensive axiomatic investigation. The present paper demonstrates that new results or extensions of well-known results can be obtained within the general framework of infinite games.
3. *Modification of BI beyond subgame perfection:* The final question is whether backward induction can be modified for other solution concepts beyond subgame perfection. An immediate ad hoc modification would consider only those SPEs that exclude partial equilibria with weakly dominated strategies, as it was demonstrated in Example 5. Perhaps, after a suitable modification of the main principle, backward induction-like reasoning could also produce some other refinements. On the other hand, proving that this is not the case would be an interesting finding as well.

Further refinements of backward induction could produce computational benefits similar to those obtained for subgame perfection. Backward solving is equivalent to the hierarchical concatenation of

solutions. Thus, solving a game with backward reasoning is equivalent to collecting together those independent solutions and connecting the global solution with stage-wise decision-making. Arguably, this is how all decisions are made.

Funding: This research received no external funding.

Acknowledgments: Michael McBride, Michael Chwe, Drew Fudenberg, Marc Kilgour, Tom Schelling, Tom Schwartz, Harrie de Swart, Donald Woodward, Alexei Zakharov, anonymous reviewers and the participants of the IMBS seminar at the University of California, Irvine, kindly provided comments.

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

Appendix A

Proof of Lemma 1. The Lemma will be proved for game G . The assumptions of finite support and crossing apply to the subgames as well as to G , which means that all steps of the proof can be repeated for every subgame of G .

Part (a): The construction of a terminal path $e \in T_\beta$ is by induction. Both when $\tau \in T_0$ and when $\tau \in T_i$ for $i = 1, \dots, n$, there is a node $v \in SU(\tau)$ such that $p_\beta^m(v) > 0$. Let us choose v as the second (after τ) node of the path.

Let us assume that a path e of length l reached a node x . If $x \in T_E$, e is the desired path. If $x \in T_D$, let I_i^k be such that $x \in I_i^k$. Both when $I_i^k \subset T_0$ and when $I_i^k \subset T_i$ for $i = 1, \dots, n$, there is a node $y \in SU(x)$ such that $p_\beta^m(y) > 0$. Path $e_y \subset \beta$ and its length is $l + 1$. The construction either ends at some endnode or proceeds indefinitely, producing some infinite path. In both cases, the resulting path e is terminal. By definition of the game and under the assumption of finite crossing, only a finite number of factors in the product $\prod_{z \in e} p_\beta^m(z)$ is in the interval $(0, 1)$; by construction, the remaining factors are equal to 1. Thus, $p_\beta(e) > 0$ and $e \in T_\beta$.

Now, let us assume that T_β is infinite. We will construct by induction an infinite path $e^* \in T_\beta$ that violates finite crossing. Path e^* begins with τ . Now, let z be such that (1) e_z crosses $k \geq 0$ information sets with nondegenerate probability distributions and (2) $\{e \in T_\beta : z \in e\}$ is infinite. We will extend e^* in such a way that the number k from (1) increases and (2) is preserved. By finite support, only for a finite number of $v \in IS(z)$, $p_\beta^m(v) > 0$; among them, for at least one t , $\{e \in T_\beta : t \in e\}$ is infinite. Let us add t to e^* . By construction, e_t crosses $k + 1$ information sets with nondegenerate probability distributions. Extending this construction indefinitely, we find a path e^* that violates finite crossing.

Part (b): If $e \notin T_\beta$, then by definition of T_β , for some $x \in e$, $p_\beta^m(x) = 0$, which implies $p_\beta(e) = 0$.

If $e \in T_\beta$, then, by final crossing, $0 < p_\beta^m(x) < 1$ only for a finite number of $x \in e$; for all other $y \in e$, $p_\beta^m(y) = 1$. This implies $p_\beta(e) > 0$.

Part (c): By (a), T_β is finite and non-empty. The thesis is proved by induction over k , i.e., for all terminal paths with the lengths shorter or equal to k or nonterminal paths with the length equal to k included in T_β .

First, the total probability of reaching the path of length 1 within T_β is 1 because $p_\beta^m(\tau) = 1$.

Second, let us assume that the thesis holds for all paths $f \subset e \in T_\beta$ that are no longer than k . Let us consider all such paths that are no longer than $k + 1$. The sum of the probabilities of reaching the final nodes of such paths includes two subsets of terms. The probabilities that are assigned to all paths no longer than k do not change and enter the summation in the same way as is the case for all paths no longer than k . There may be new terminal or nonterminal paths with the exact length of $k + 1$. Here, two cases are possible. First, the probability of reaching the extra node $k + 1$ is 1. This means that the term associated with such a path f enters the summation in the same way as for all paths no longer than k . Second, the distribution at the next-to-final node y is nondegenerate. In such a case, instead of a single path e that is no longer than k , we now have—by finite support—a finite number of paths $\{f_l\}_{l \in L}$ that are no longer than $k + 1$ and that go through y . The total probability associated to them at node y is 1. Thus, the summation term corresponding to path e is now substituted by a finite number of terms that add up to the same number. This means that the total probability of reaching the ends of all distinct paths that are no longer than $k + 1$ is 1. Since T_β is finite, finite crossing implies that

there is a common maximum for the length of the segments, such that the associated probabilities are smaller than 1. The probabilities associated with reaching this maximum—or a shorter length if the path ends earlier—is equal to the probability distribution over T_β and is equal to 1. \square

Proof of Lemma 2. Part (a): Every subgame H divides the set of information sets of player i into two subsets. Every β_i can be partitioned into two partial strategies, β_i^H and β_i^{-H} , defined for those two subsets the same way β_i is defined. First, two different strategies produce two distinct pairs of partial strategies, since they must differ for at least one information set. Second, every pair of partial strategies β_i^H and β_i^{-H} may result from some partition—namely, from the partition of a strategy that is identical for all information sets with the pair.

Part (b): It is a simple consequence of (a). \square

Proof of Lemma 3. A terminal path in \mathbf{s} either includes ϕ or does not include ϕ . By Lemma 1, the total number of paths in \mathbf{s} is finite; hence, $P^G(\mathbf{s})$ can be represented as the sum of the payoffs

$$P^G(\mathbf{s}) = \sum_{e \in T_{s(H)}} p_{sG}(e) P^G(e) + \sum_{e \in T_{s(G-H)}} p_{sG}(e) P^G(e) \tag{A1}$$

We need to show that the first term in Equation (A1) is equal to the payoff of \mathbf{s} in H multiplied by the probability of reaching ϕ , the root of H .

If $p_{sG}(\phi) = 0$, then, from the definition of $p_{sG}(e)$, we have the following:

$$p_{sG}(e) = 0 \text{ for all } e \in T_{s(H)}$$

Since the summation is over a finite set, this means that

$$\sum_{e \in T_{s(H)}} p_{sG}(e) P^G(e) = 0 = p_{sG}(\phi) P^H(\mathbf{s}^H).$$

Now, let us assume that $p_{sG}(\phi) > 0$. First, notice that

- (i) every terminal path $e \in T_{s(H)}$ defines a terminal path of \mathbf{s}^H in H and all terminal paths of \mathbf{s}^H in H can be obtained this way.

Moreover, for every $e \in T_{s(H)}$ and the corresponding $e^H \in T_{s^H}^H$, we have the following:

- (ii) $P^G(e) = P^H(e^H)$ and

$$(iii) p_{sG}(e) = \prod_{y \in e} p_s^m(y) = \prod_{y \in e^H - \{\phi\}} p_s^m(y) \prod_{y \in e - e^H \cup \{\phi\}} p_s^m(y) = p_{sG}(\phi) \times p_{s^H}(e^H)$$

We can then use (i)–(iii) to make substitutions:

$$\sum_{e \in T_{s(H)}} p_{sG}(e) P^G(e) = \sum_{e^H \in T_{s^H}^H} p_{sG}(\phi) \times p_{s^H}(e^H) \times P^H(e^H) = p_{sG}(\phi) P^H(\mathbf{s}^H).$$

\square

Proof of Lemma 4. By Lemma 3, it is sufficient to show that

$$P^F(\mathbf{s}^{-H}) = p_{sG}(\phi) P^H(e_\phi) + \sum_{e \in T_{s(G-H)}} p_{sG}(e) P^G(e) \tag{A2}$$

This follows from the fact that \mathbf{s}^{-H} is identical to \mathbf{s} for all information sets outside of H and that all terminal paths for \mathbf{s}^{-H} in F , with the exception of e_ϕ , have by definition the same probabilities and payoffs assigned as in G . For e_ϕ , the probabilities of reaching ϕ are equal by definition; the equality $P^H(e_\phi) = P^F(\phi)$ follows from the definition of F as an (\mathbf{s}, H) -upgame of G . \square

Proof of Theorem 1. (a) \rightarrow (b): decomposing an SPE strategy profile must result in a pair of SPE strategy profiles.

Since H is a subgame of G and since \mathbf{s}^G is an SPE in G , \mathbf{s}^H must be an SPE for H . We need to prove that \mathbf{s}^F is an SPE for F . Let us assume that this is not the case and that we can find J —a subgame of F —such that \mathbf{s}^J is not a Nash equilibrium. Here, two cases are possible:

Case 1: J does not include ϕ . In this case, J is disjoint with H . Thus, J is also a subgame of G and \mathbf{s}^J is identical with \mathbf{s}^{FJ} . However, (a) implies that \mathbf{s}^J is a Nash equilibrium, which also must be the case for \mathbf{s}^{FJ} .

Case 2: J includes ϕ . Let us assume that there is a player, i , of which the strategy t_i^J gives a higher payoff than $P_i(\mathbf{s}^J)$, i.e.,

$$P_i^J(\mathbf{s}_{-i}^J, t_i^J) > P_i^J(\mathbf{s}^J) \tag{A3}$$

Let us denote $(\mathbf{s}_{-i}^J, t_i^J)$ as \mathbf{t}^J . We will find K , a subgame of G , such that all payoffs in K are identical with the corresponding payoffs in J . Let K be defined as J , with ϕ substituted with subgame H . By construction, K is a subgame of G and \mathbf{s}^K is a Nash equilibrium in K . J is a (\mathbf{s}, H) -upgame of K . Let us define a new strategy profile in K as $\mathbf{t}^K = \mathbf{t}^J \cup \mathbf{s}^H$. We now apply Lemma 4 to K ; J ; strategy profiles $\mathbf{t}^J, \mathbf{t}^K, \mathbf{s}^J$, and \mathbf{s}^K ; and player i :

$$P_i^K(\mathbf{t}^K) = P_i^J(\mathbf{t}^J) \tag{A4}$$

and

$$P_i^K(\mathbf{s}^K) = P_i^J(\mathbf{s}^J) \tag{A5}$$

Subtracting the respective sides of Equation (A5) from Equation (A4), we obtain a contradiction with our assumption that \mathbf{s}^K is an NE in K :

$$P_i^K(\mathbf{t}^K) - P_i^K(\mathbf{s}^K) = P_i^J(\mathbf{t}^J) - P_i^J(\mathbf{s}^J) > 0 \tag{A6}$$

(b) \rightarrow (a): every strategy profile resulting from the concatenation of SPE strategy profiles is SPE.

Let \mathbf{s}^H be an SPE in H , the subgame of G ; \mathbf{s}^F be SPE in F , the (\mathbf{s}, H) -upgame of G ; and let $\mathbf{s} = \mathbf{s}^H \cup \mathbf{s}^F$. We need to show that \mathbf{s} is SPE.

Let us assume that this is not the case. Thus, we can find a subgame K of G such that \mathbf{s}^K is not a Nash equilibrium. This implies that some player i could improve their payoff against \mathbf{s}_{-i}^K by playing some strategy t_i^K , i.e.,

$$P_i^K(\mathbf{s}_{-i}^K, t_i^K) > P_i^K(\mathbf{s}^K) \tag{A7}$$

I will show that Equation (A7) is in contradiction with (b). Let us denote $(\mathbf{s}_{-i}^K, t_i^K)$ as (\mathbf{t}^K) . Three cases are possible.

Case 1: K is a subgame of F that does not include ϕ . By the subgame perfection of F and contrary to Equation (A7), \mathbf{s}^K must be a Nash equilibrium.

Case 2: K is a subgame of H . By the subgame perfection of H and contrary to Equation (A7), \mathbf{s}^K must be a Nash equilibrium.

Case 3: K includes the node ϕ and at least one more node from F . In such a case, K must include all nodes that follow ϕ and H must be a subgame of K . Let J be the (\mathbf{s}^K, H) -upgame of K . J is also a subgame of F . By (b), F is an SPE and both \mathbf{s}^{FH} and $\mathbf{s}^{FJ} = \mathbf{s}^J$ are SPEs in H and J , respectively. By Lemma 4,

$$P_i^K(\mathbf{t}^K) = P_i^J(\mathbf{t}^J) \tag{A8}$$

and

$$P_i^K(\mathbf{s}^K) = P_i^J(\mathbf{s}^J) \tag{A9}$$

Subtracting the respective sides of Equation (A9) from Equation (A8), we obtain a contradiction with our assumptions:

$$P_i^J(\mathbf{t}^J) - P_i^J(\mathbf{s}^J) = P_i^K(\mathbf{t}^K) - P_i^K(\mathbf{s}^K) > 0 \tag{A10}$$

□

Proof of Lemma 5. First, let us assume that neither the root ϕ of H follows ψ —the root of J —nor ψ follows ϕ . If there is a third node, χ , that belongs to both subgames, then, by definition of the subgame,

both ϕ and ψ must be followed by χ . Thus, we could find at least two different paths to χ : one through ϕ and one through ψ , which is inconsistent with the definition of the tree.

Now, let us assume that ϕ follows ψ or vice versa. However, this would imply that the path to ϕ in the agenda is longer than the path to ψ (or vice versa) and that the subgames cannot be at the same level. \square

Proof of Theorem 2. Subgame perfection of H_θ is assumed on both sides, so we need to prove that \mathbf{s} is SPE for G iff \mathbf{s}^F is SPE for F .

Part 1: Θ is finite, i.e., $\Theta = \{1, \dots, n\}$. Theorem 2 follows from applying Theorem 1 in both directions a finite number of times to a chain of upgames $\{G_i\}_{i=1, \dots, n+1}$ such that $G = G_1$, $G_{n+1} = F$, and, for $i = 1, \dots, n$, G_{i+1} is the $(\mathbf{s}^{G_i}, \theta_i)$ -upgame of G_i .

Part 2: Θ is infinite.

(a) \rightarrow (b):

Let us check if it is possible to find a strategy t_i^F in F for some player i that would offer a higher payoff in some subgame L of F than \mathbf{s} . Let us denote the strategy profile $(\mathbf{s}_{-i}^F, t_i^F)$ as \mathbf{t}^F . By Lemma 1, the total number of terminal paths relevant for \mathbf{s}^F , \mathbf{t}^F , or \mathbf{t}^L is finite. Since all subgames with their roots in Θ are disjoint, this implies that the number of roots in Θ that are part of any such path is also finite. Let us denote the set of such roots as Δ and the (\mathbf{s}, Δ) -upgame of G as K . Since Δ is finite, \mathbf{s}^K is SPE in K by the finite part of the proof (Part 1).

Let us denote as M the subgame of K that has the same root as L . By construction of K , $M=L$ or L is an (\mathbf{s}^M, Δ) -upgame of M . Since \mathbf{s}^K is SPE, \mathbf{s}^M is SPE, and by Part 1 of the Proof, \mathbf{s}^L is SPE; player i cannot receive a higher payoff with t_i^L .

(b) \rightarrow (a): The proof is similar to the (a) \rightarrow (b) part. The assumption that there exists a strategy t_i that could improve the payoff of player i in some subgame L of G leads to a contradiction. The strategy profiles \mathbf{s} , $(\mathbf{s}_{-i}, t_i) = \mathbf{t}$, and \mathbf{t}^L have only a finite number of paths that reach at least one of $\{H_\theta\}_{\theta \in \Theta}$. This allows us to drop an infinite number of subgames from Θ and to construct a smaller subgame K of L with the payoffs identical to L for respective strategy profiles \mathbf{t}^L , \mathbf{s}^L , \mathbf{t}^K , and \mathbf{s}^K such that F is an (\mathbf{s}^K, Δ) -upgame of K for some finite subset $\Delta \subset \Theta$. By Part 1, \mathbf{s}^K is SPE for K , which means that an improvement in payoff for player i playing s_i^K is impossible in K , which implies that it is also impossible for i to improve from playing s_i^L in L . \square

References

1. Zermelo, E. Über Eine Anwendung Der Mengenlehre Auf Die Theorie Des Schachspiels. In Proceedings of the Fifth International Congress of Mathematicians, Cambridge, UK, 21–28 August 1912; Cambridge University Press: Cambridge, UK, 1913; Volume 2, pp. 501–504.
2. Schwalbe, U.; Walker, P. Zermelo and the Early History of Game Theory. *Games Econ. Behav.* **2001**, *34*, 123–137. [[CrossRef](#)]
3. Von Stackelberg, H. *Market Structure and Equilibrium (Marktform Und Gleichgewicht)*; Springer: New York, NY, USA, 2011.
4. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior*; Princeton University Press: Princeton, NJ, USA, 1944; ISBN 978-0-691-13061-3.
5. Von Neumann, J. Zur Theorie Der Gesellschaftsspiele. *Math. Ann.* **1928**, *100*, 295–320. [[CrossRef](#)]
6. Kuhn, H. Extensive Games and the Problem of Information. In *Contributions to the Theory of Games I*; Kuhn, H., Tucker, A., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; pp. 193–216.
7. Schelling, T.C. *The Strategy of Conflict*; Harvard University Press: Cambridge, MA, USA, 1960.
8. Selten, R. Spieltheoretische Behandlung Eines Oligopolmodells Mit Nachfrageträgheit: Teil i: Bestimmung Des Dynamischen Preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft. J. Inst. Theor. Econ.* **1965**, *2*, 301–324.
9. Rosenthal, R.W. Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox. *J. Econ. Theory* **1981**, *25*, 92–100. [[CrossRef](#)]
10. Selten, R. The Chain Store Paradox. *Theory Decis.* **1978**, *9*, 127–159. [[CrossRef](#)]

11. Selten, R. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *Int. J. Game Theory* **1975**, *4*, 25–55. [\[CrossRef\]](#)
12. Kreps, D.M.; Wilson, R. Sequential Equilibria. *Econometrica* **1982**, *50*, 863–894. [\[CrossRef\]](#)
13. Kohlberg, E.; Mertens, J.-F. On the Strategic Stability of Equilibria. *Econometrica* **1986**, *54*, 1003–1037. [\[CrossRef\]](#)
14. Basu, K. Strategic Irrationality in Extensive Games. *Math. Soc. Sci.* **1988**, *15*, 247–260. [\[CrossRef\]](#)
15. Basu, K. On the Non-Existence of a Rationality Definition for Extensive Games. *Int. J. Game Theory* **1990**, *19*, 33–44. [\[CrossRef\]](#)
16. Bicchieri, C. Self-Refuting Theories of Strategic Interaction: A Paradox of Common Knowledge. In *Philosophy of Economics*; Springer: New York, NY, USA, 1989; pp. 69–85. [\[CrossRef\]](#)
17. Binmore, K. Modeling Rational Players: Part I. *Econ. Philos.* **1987**, *3*, 179–214. [\[CrossRef\]](#)
18. Binmore, K. Modeling Rational Players: Part II. *Econ. Philos.* **1988**, *4*, 9–55. [\[CrossRef\]](#)
19. Bonanno, G. The Logic of Rational Play in Games of Perfect Information. *Econ. Philos.* **1991**, *7*, 37–65. [\[CrossRef\]](#)
20. Fudenberg, D.; Kreps, D.M.; Levine, D.K. On the Robustness of Equilibrium Refinements. *J. Econ. Theory* **1988**, *44*, 354–380. [\[CrossRef\]](#)
21. Pettit, P.; Sugden, R. The Backward Induction Paradox. *J. Philos.* **1989**, *86*, 169–182. [\[CrossRef\]](#)
22. Reny, P. *Rationality, Common Knowledge and the Theory of Games*; Princeton University Press: Princeton, NJ, USA, 1988.
23. Aliprantis, C.D. On the Backward Induction Method. *Econ. Lett.* **1999**, *64*, 125–131. [\[CrossRef\]](#)
24. Schelling, T.C. (University of Maryland, College Park, MD, USA). Personal communication, 2008.
25. Fudenberg, D.; Tirole, J. *Game Theory*; MIT Press: Cambridge, MA, USA, 1991.
26. Myerson, R.B. *Game Theory*; Harvard University Press: Cambridge, MA, USA, 2013.
27. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; MIT Press: Cambridge, MA, USA, 1994; ISBN 0-262-14041-7.
28. Escardò, M.H.; Oliva, P. Computing Nash Equilibria of Unbounded Games. *Turing-100* **2012**, *10*, 53–65.
29. Neyman, A.; Sorin, S. *Stochastic Games and Applications*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 2003; ISBN 1-4020-1492-9.
30. Aumann, R.J. *Mixed and Behavior Strategies in Infinite Extensive Games*; Report; Princeton University: Princeton, NJ, USA, 1961.
31. Nash, J. Non-Cooperative Games. *Ann. Math.* **1951**, *54*, 286–295. [\[CrossRef\]](#)
32. Bendor, J.; Swistak, P. Evolutionary Equilibria: Characterization Theorems and Their Implications. *Theory Decis.* **1998**, *45*, 99–159. [\[CrossRef\]](#)
33. Flood, M.M. *Some Experimental Games*; Research Memorandum RM-789; RAND: Santa Monica, CA, USA, 1952.
34. Brams, S.J.; Taylor, A.D. *Fair Division: From Cake-Cutting to Dispute Resolution*; Cambridge University Press: Cambridge, UK, 1996; ISBN 0-521-55390-3.
35. Steinhaus, H. The Problem of Fair Division. *Econometrica* **1948**, *16*, 101–104.
36. Romer, T.; Rosenthal, H. Political Resource Allocation, Controlled Agendas, and the Status Quo. *Public Choice* **1978**, *33*, 27–43. [\[CrossRef\]](#)
37. Kaminski, M.M.; Nalepa, M. A Model of Strategic Preemption: Why Do Post-Communists Hurt Themselves? *Decisions* **2014**, *21*, 31–65. [\[CrossRef\]](#)
38. Kaminski, M.M.; Lissowski, G.; Swistak, P. The “Revival of Communism” or the Effect of Institutions? The 1993 Polish Parliamentary Elections. *Public Choice* **1998**, *97*, 429–449. [\[CrossRef\]](#)
39. Farquharson, R. *Theory of Voting*; Blackwell: Hoboken, NJ, USA, 1969.
40. Riker, W.H. *The Art of Political Manipulation*; Yale University Press: New Haven, CT, USA, 1986; ISBN 0-300-03592-6.

