

Milchtaich, Igal

Working Paper

Schedulers, potentials and weak potentials in weakly acyclic games

Working Paper, No. 2013-03

Provided in Cooperation with:

Department of Economics, Bar-Ilan University

Suggested Citation: Milchtaich, Igal (2013) : Schedulers, potentials and weak potentials in weakly acyclic games, Working Paper, No. 2013-03, Bar-Ilan University, Department of Economics, Ramat-Gan

This Version is available at:

<https://hdl.handle.net/10419/96060>

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.

Schedulers, Potentials and Weak Potentials in Weakly Acyclic Games

Igal Milchtaich*

May 2013

Abstract. In a number of large, important families of finite games, not only do pure-strategy Nash equilibria always exist but they are also reachable from any initial strategy profile by some sequence of myopic single-player moves to a better or best-response strategy. This *weak acyclicity* property is shared, for example, by all perfect-information extensive-form games, which are generally not *acyclic* since even sequences of best-improvement steps may cycle. Weak acyclicity is equivalent to the existence of *weak potential*, which unlike a *potential* increases along some rather than every sequence as above. It is also equivalent to the existence of an acyclic *scheduler*, which guarantees convergence to equilibrium by disallowing certain (improvement) moves. A number of sufficient conditions for acyclicity and weak acyclicity are known.

Keywords. Weakly acyclic games, Weak potential, Scheduler.

This paper concerns *finite* games, with a finite number n of players and a finite strategy set S_i for each player i . Correspondingly, “strategy” always means pure strategy. The payoff function of player i is denoted by h_i . A *subgame* of a finite game Γ is obtained by replacing each strategy set S_i with some subset of S_i and restricting the payoff functions correspondingly.¹ If the strategy sets of one or more players are reduced to a singleton, it is possible to view only the remaining ones as players in the subgame.

The *improvement graph* of a finite game Γ is the directed graph that describes the players’ profitable unilateral deviations. Its vertices are the strategy profile in the game, and for every pair of strategy profiles s and t , a (directed) edge with head s and tail t exists if and only if there is some player i such that $s_j = t_j$ for all $j \neq i$ (thus, $s = (s_i, t_{-i})$) and

$$h_i(s) > h_i(t). \tag{1}$$

The *best-improvement graph* of Γ is the subgraph obtained by augmenting (1) with the requirement that, for player i , strategy s_i is a best response to $s_{-i} (= t_{-i})$, that is, $h_i(s) \geq h_i(s'_i, s_{-i})$ for all other strategies s'_i . Obviously, a strategy profile is a sink of either the improvement or best-improvement graph if and only if it is a (pure-strategy Nash) *equilibrium* in Γ . In the following, “the graph” of Γ and related terms may refer to either

*Department of Economics, Bar-Ilan University, Ramat Gan 5290002, Israel
igal.milchtaich@biu.ac.il <http://faculty.biu.ac.il/~milchti>

¹ It should be clear from the context whether “subgame” is meant in this sense (Shapley 1964) or in the more familiar one pertaining to extensive-form games (Selten 1975). For example, the latter holds when the reference is to an extensive-form game G , and the former holds when it is to G ’s normal, or strategic, form.

graph. This facilitates the simultaneous presentation of two parallel terminologies. However, unless this can be understood from the context, unambiguous use of any term requires indicating the graph it refers to, for example, by prefixing either *I*- or *BI*-.²

A *scheduler*³ for a finite game Γ is any subgraph of its graph that includes all vertices and at least one edge with tail t for every vertex t that has such edges (i.e., is not a sink). Less formally, a scheduler is a rule that, for each strategy profile, may restrict the players' freedom of choice by not allowing some of them to take certain moves, or even any move at all, without going as far as making it impossible to leave the strategy profile. For example, it may stipulate that certain kinds of moves take precedence over certain other moves, so that, whenever any of the former is feasible, none of the latter is allowed. One scheduler is *weaker* than another if it is a subgraph of it. The second, *stronger* scheduler allows every move allowed by the first one but not necessarily the other way around.⁴ The strongest scheduler, which is the (improvement or best-improvement) graph itself, is also referred to as the *default scheduler*.

A finite sequence s^0, s^1, \dots, s^m ($m \geq 0$) of (not necessarily distinct) strategy profiles is a *walk* of length m in a scheduler if for $l = 1, 2, \dots, m$ there is an edge in the scheduler whose tail and head are s^{l-1} and s^l , respectively. A walk is *closed* if $m > 0$ and $s^m = s^0$, and it is a *path* if the $m + 1$ strategy profiles are all distinct. One walk or path *extends* another if the former is obtained from the latter by appending to it one or more strategy profiles. A scheduler is *acyclic* if there are no closed walks in it, and *weakly acyclic* if some weaker scheduler is acyclic.

The game itself is said to be acyclic or weakly acyclic if the default scheduler has the same property. A path in the default scheduler is also called an *improvement* or *best-(reply) improvement* path, depending on the graph considered. Correspondingly, alternative terms for the (weak) *I*- and *BI*-acyclicity properties are the (respectively, *weak*) *finite improvement* and *finite best-(reply) improvement properties*.⁵ It is easy to see that the four properties of games are linearly ordered by the implication relation, as follows:

$$I\text{-acyclicity} \Rightarrow BI\text{-acyclicity} \Rightarrow \text{weak } BI\text{-acyclicity} \Rightarrow \text{weak } I\text{-acyclicity}.$$

² A third graph considered in the literature is the *best-reply* graph, which differs from the best-improvement one in that it describes also moves between (and not only to) best-response strategies. A strategy profile is a sink of the best-reply graph if and only if it is a *strict* equilibrium. The set of edges in the best-improvement graph is the intersection of those in the best-reply and the improvement graphs.

³ The meaning of this term here is somewhat different than in Apt and Simon (2012). According to these authors' definition, at each strategy profile, a scheduler allows only one player to move but does not restrict his choice of strategy. In addition, the identity of the mover may depend on history, that is, on previous moves.

⁴ Note that these definitions entail reflexivity: every scheduler is both weaker and stronger than itself. This fact may optionally be underlined by adding the qualifier *weakly*. The corresponding irreflexive relation is indicated by the qualifier *strictly*.

⁵ Young's (1993) notion of (weak) acyclicity is similar, except that it refers to the best-reply graph (see footnote 2), and may therefore be referred to as (respectively, weak) *BR-acyclicity*. *BR-acyclicity* in particular precludes the existence of *best-response cycles* in the sense of Voorneveld (2000). The latter differ from closed walks in the best-improvement graph of the game in that only *one* of the changes of strategy is required to be an improvement; the rest may be moves between two best-response strategies.

For some examples of games possessing one or more of these properties see Monderer and Shapley (1996), Milchtaich (1996, 2009), Friedman and Mezzetti (2001), Milchtaich and Winter (2002), Kukushkin et al. (2005), Engelberg and Schapira (2011) and Theorems 2, 3, 4 and 5 below.

A real-valued function P on the set of vertices is a *potential* for a scheduler if it increases along every walk in it, in other words, if for every two strategy profiles s and t that are respectively the head and tail of an edge in the scheduler,

$$P(s) > P(t). \quad (2)$$

A function P is a *weak potential* for a scheduler if it is a potential for some weaker scheduler, equivalently, if the subgraph obtained by eliminating all edges in the scheduler whose head and tail do not satisfy (2) is also a scheduler. A necessary and sufficient condition for this is that every strategy profile t that is not an equilibrium is also not a “local minimum point” of P , in the sense that t is the tail of some edge in the scheduler whose head s satisfies (2).

A potential or weak potential for a *game* means such a function for the default scheduler. An alternative term for I -potential for a game, which stresses the distinction between this concept and the related cardinal one of exact potential, is *generalized ordinal potential* (Monderer and Shapley 1996). It is easy to see that the following implications between properties of a function P on strategy profiles hold:

$$I\text{-potential} \Rightarrow BI\text{-potential} \Rightarrow \text{weak } BI\text{-potential} \Rightarrow \text{weak } I\text{-potential}.$$

The following theorem applies to both the improvement and best-improvement graph.

Theorem 1. (Monderer and Shapley 1996, Kukushkin 2004) For a finite game or (more generally) a scheduler, the following properties are equivalent:

- (i) acyclicity,
- (ii) existence of potential,
- (iii) every walk can be extended only finitely many times before an equilibrium is reached.

Similarly, the following properties are equivalent:

- (i') weak acyclicity,
- (ii') existence of weak potential,
- (iii') for every strategy profile s , *some* path that starts at s ends at an equilibrium.

Proof. For an acyclic scheduler (or, as a special case, acyclic game), consider for each strategy profile s the length of the longest path that starts at s . This number is 0 if and only if s is an equilibrium. Its negative,

$$P(s) = -\max\{m \geq 0 \mid \text{there is a path of length } m \text{ that starts at } s\},$$

defines a potential, as it is easy to see that P increases along any walk. Conversely, if a scheduler does have a closed walk, the latter clearly precludes the existence of a potential, and it can be extended indefinitely by repetition.




Figure 1. The cube-like improvement (and best-improvement) graph of a finite $2 \times 2 \times 2$ game that is (i) weakly acyclic but (ii) not acyclic, (iii) not solvable by iterated elimination of never-best response strategies and (iv) has the property that every subgame (which corresponds to a k -dimensional face of the cube, with $0 \leq k \leq 3$) has a unique equilibrium. The strategies are right and left for player 1, up and down for player 2, and in and out for player 3. Shown for each strategy profile are the players' payoff vector and the value of a weak potential P defined as in the proof of Theorem 1.

By the first part of the proof, a scheduler is weakly acyclic if and only if some weaker scheduler possesses a potential, in other words, if and only if the scheduler itself possesses a weak potential. In this case, every walk in the weaker scheduler that starts at a given strategy profile s is a path and can be extended only finitely many times before it reaches an equilibrium, which proves that (iii') holds. Conversely, for a scheduler that satisfies (iii'), consider for each strategy profile s the distance in the scheduler to the closest equilibrium. The negative of this distance (see example in Figure 1),

$$P(s) = -\min\{m \geq 0 \mid \text{there is a path of length } m \text{ that starts at } s \text{ and ends in an equilibrium}\},$$

defines a function P on strategy profiles that is a weak potential for the scheduler. This is because, if a strategy profile s is not an equilibrium, then P increases along any of the shortest paths connecting it to an equilibrium, which in particular means that s is not a "local minimum point" (see above). ■

Sufficient Conditions for Weak Acyclicity

In a finite game Γ , a *never-best response* strategy for a player i is a strategy s_i that is not a best response to any profile s_{-i} of the other players' strategies. Since, as indicated, only pure strategies are considered here, every strategy that is strictly dominated by a *mixed* strategy is a never-best response strategy but not conversely. *Iterated elimination of never-best response strategies* means a finite sequence $\Gamma^0, \Gamma^1, \dots, \Gamma^m$ ($m \geq 0$) of games such that $\Gamma^0 = \Gamma$, each of the subsequent games is a subgame of the preceding game obtained by eliminating one or more never-best response strategies for one or more players, and there are no never-best response strategies in Γ^m .⁶ As the following lemma shows, the elimination process in a sense preserves the best-response relation.

⁶ It can be shown that the order of elimination does not matter, in the sense that the last subgame Γ^m is unique. Obviously, the other subgames are also unique if in each step *all* eligible strategies are eliminated.

Lemma 1. Consider iterated elimination of never-best response strategies in a finite game Γ , that is, a sequence of subgames $\Gamma^0 (= \Gamma), \Gamma^1, \dots, \Gamma^m$ as above. For $0 \leq l \leq m$, and any strategy profile s in Γ^l and player i , strategy s_i is a best response to s_{-i} for player i in Γ^l if and only if this is so in Γ . Moreover, for $0 \leq l \leq m$, the set of all equilibria in Γ^l coincides with that in Γ .

Proof. For $0 \leq l \leq m$, and any strategy profile s in Γ^l and player i , if s_i is *not* a best response to s_{-i} for player i in Γ^l , then the same is obviously true in Γ . Conversely, if s_i is not a best response to s_{-i} in Γ , then consider any strategy s'_i that is a best response. Since the games $\Gamma^0, \Gamma^1, \dots, \Gamma^l$ all include every strategy in s_{-i} , the same must be true for the best response strategy s'_i , which does not get eliminated. This implies that strategy s_i is not a best response also in Γ^l . The second part of the lemma follows from the first part and the fact that, for $0 < l \leq m$, every equilibrium in Γ^{l-1} is present also in Γ^l , since each of the strategies in it is a best response to the others. ■

A game Γ is *solvable* by iterated elimination of never-best response strategies if there exists a sequence as above such that, in the last subgame Γ^m , all strategy profiles are equilibria. By Lemma 1, in this case, the set of all strategy profiles in Γ^m is also the set of all equilibria in Γ . In other words, solvability means that the equilibria in Γ are the only strategy profile that survive iterated elimination of never-best response strategies.

Theorem 2. (Kukushkin 2012, Apt and Simon 2012) If a finite game is solvable by iterated elimination of never-best response strategies, then it is weakly *BI*-acyclic.

Proof. Consider a game Γ solvable by iterated elimination of never-best response strategies and a corresponding sequence of subgames $\Gamma^0, \Gamma^1, \dots, \Gamma^m$. Define the *height* of a strategy in Γ as the largest index l such that Γ^l includes the strategy, and the height of a strategy profile as the average height of the strategies in it. It suffices to establish the following.

CLAIM. The function P that maps strategy profile to their height is a weak potential for Γ . A strategy profile s is an equilibrium if and only if $P(s) = m$.

If all the strategies in a strategy profile s are in Γ^m , then by the solvability assumption s is an equilibrium. Otherwise, consider a strategy s_i in s with minimum height $k (< m)$, so that all the strategies in s are in Γ^k but (at least) s_i is not in Γ^{k+1} . Necessarily, for the corresponding player i , strategy s_i is not a best response to s_{-i} (in both Γ^k and Γ ; see Lemma 1). Let s'_i be a strategy in Γ^k that is a best response. The height of s'_i is greater than k , and therefore the strategy profile (s'_i, s_{-i}) satisfies $P(s) < P(s'_i, s_{-i}) \leq m$. This concludes the proof of the claim, and hence also that of the theorem. ■

A different sufficient condition for weak acyclicity is presented by the following theorem. See Figure 1 for example.

Theorem 3. (Fabrikant et al. 2010) If a finite game has the property that every subgame has a unique equilibrium, then it is weakly *BI*-acyclic.

Interestingly, the weaker property that every subgame has *at least* one equilibrium is not sufficient even for weak *I*-acyclicity (Takahashi and Yamamori 2002).




Figure 2. An extensive-form game with perfect information that is not *BI*-acyclic. The best-improvement graph of the game has the following closed walk: (L, l'') , (R, l'') , (R, r'') , (L, r'') and back to (L, l'') .

Extensive-Form Games

The difference between acyclicity and weak acyclicity is illustrated by the example of finite extensive-form games with perfect information. In general, these games are not even *BI*-acyclic, as the simple game in Figure 2 demonstrates. However, that particular counterexample is clearly driven by player 2's nonbeneficial change of action at his unreached decision node (simultaneously with the beneficial one at the reached node). Indeed, in the (trivial) subgame whose root is that node, the change of strategy is harmful rather than beneficial. In other words, the agent (Selten 1975) residing at the node switches to an action that would not be optimal if the node were actually reached. The significance of this observation lies in the fact that the agent normal form of every finite extensive-form game with perfect information *is* acyclic. This fact can be stated also as follows.

Theorem 4. (Kukushkin 1999) Every finite extensive-form game with perfect information in which each player has only one decision node is *I*-acyclic.

Theorem 4 is a special case of a more general result, which applies regardless of the numbers of decision nodes. Namely, if a walk in the improvement graph does not involve changes of actions at unreached decision nodes as in Figure 2, then it cannot be closed (Kukushkin 2002). In other words, the scheduler that results from forbidding such changes of actions is acyclic. An alternative (similar, but a trifle weaker) acyclic scheduler can be defined as follows. For a player i and a strategy profile s , call a unilateral change of strategy by i from s_i to some other strategy s'_i *parsimonious* if $h_i(s''_i, s_{-i}) < h_i(s'_i, s_{-i})$ for every strategy s''_i that differs from both s_i and s'_i but is a combination of them, in the sense that the action it prescribes at each of player i 's decision nodes is also prescribed by one of these strategies. Clearly, any non-parsimonious change of strategy can be replaced by a parsimonious one without decreasing the resulting payoff of the player involved. Therefore, the subgraphs of the improvement- and best-improvement graphs obtained by considering only parsimonious changes of strategies are schedulers.

Lemma 2. In every finite extensive-form game with perfect information G , the *I*-scheduler and *BI*-scheduler defined by parsimonious changes of strategies are acyclic.

Proof. It needs to be shown that any walk in the I -scheduler (or, as a special case, the BI -scheduler) under consideration, s^0, s^1, \dots, s^m , is not closed. For a subgame \hat{G} of G and a strategy profile s , denote by \hat{s} the strategy profile in \hat{G} obtained by restricting each player's strategy to the decision nodes in that subgame. In particular, $\hat{s}^0, \hat{s}^1, \dots, \hat{s}^m$ are the strategy profiles corresponding to the walk. Now, choose \hat{G} in such a way that there is exactly one player i whose strategies $\hat{s}_i^0, \hat{s}_i^1, \dots, \hat{s}_i^m$ are not all equal. For any $0 \leq l < m$ such that $\hat{s}_i^l \neq \hat{s}_i^{l+1}$, consider the strategy $s_i^{l+1/2}$ ($\neq s_i^{l+1}$) in G that coincides with s_i^l inside \hat{G} and with s_i^{l+1} outside it. Since the chance from s_i^l to s_i^{l+1} increases i 's payoff and is parsimonious, $h_i(s_i^{l+1/2}, s_{-i}^l) < h_i(s_i^{l+1}, s_{-i}^l)$. The inequality implies that, in the subgame \hat{G} , strategy \hat{s}_i^{l+1} yields player i a higher payoff than \hat{s}_i^l against the other players' strategies \hat{s}_{-i}^l . Since, by assumption, the latter do not change (that is, $\hat{s}_{-i}^0 = \hat{s}_{-i}^1 = \dots = \hat{s}_{-i}^m$), this conclusion implies that the walk cannot be closed. ■

The guaranteed existence of an acyclic scheduler means that, while general perfect-information extensive-form games may not be acyclic, they are always weakly acyclic.

Theorem 5. (Kukushkin 2002) Every finite extensive-form game with perfect information is weakly BI -acyclic.

By Theorems 4 and 5, acyclicity or weak acyclicity of a finite game Γ is a necessary condition for the existence of *some* perfect-information extensive-form game G whose normal or agent normal form, respectively, is Γ . If Γ is the agent normal form of G , then it has the additional property that every subgame has an equilibrium. However, if Γ is the normal form, this is not necessarily so. For example, if in Figure 2 player 2 were only allowed to use strategies $r'r''$ and $l'l''$, an equilibrium would not exist.

References

- Apt, K. R. and Simon, S. (2012). A classification of weakly acyclic games. Lecture Notes in Computer Science 7615, 1–12.
- Engelberg, R. and Schapira, M. (2011) Weakly-acyclic (internet) routing games. Lecture Notes in Computer Science 6982, 290–301.
- Fabrikant, A., Jaggard, A. D. and Schapira, M. (2010). On the structure of weakly acyclic games. Lecture Notes in Computer Science 6386, 126–137.
- Friedman, J. W. and Mezzetti, C. (2001). Learning in games by random sampling. Journal of Economic Theory 98, 55–84.
- Kukushkin, N. S. (1999). Potential games: a purely ordinal approach. Economics Letters 64, 279–283.
- Kukushkin, N. S. (2002). Perfect information and potential games. Games and Economic Behavior 38, 306–317.

- Kukushkin, N. S. (2004). Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior* 48, 94–110.
- Kukushkin, N. S. (2012). Cournot tâtonnement and dominance solvability in finite games. *Discrete Applied Mathematics* 160, 948–958.
- Kukushkin, N. S., Takahashi, S. and Yamamori, T. (2005). Improvement dynamics in games with strategic complementarities. *International Journal of Game Theory* 33, 229–238.
- Milchtaich, I. (1996). [Congestion games with player-specific payoff functions](#). *Games and Economic Behavior* 13, 111–124.
- Milchtaich, I. (2009). [Weighted congestion games with separable preferences](#). *Games and Economic Behavior* 67, 750–757.
- Milchtaich, I. and Winter, E. (2002). [Stability and segregation in group formation](#). *Games and Economic Behavior* 38, 318–346.
- Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior* 14, 124–143.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 25–55.
- Shapley, L. S. (1964). Some topics in two-person games. In: *Advances in Game Theory (Annals of Mathematics Studies 52)*. M. Dresher, L. S. Shapley, and A. W. Tucker (Eds.), Princeton: Princeton University Press, 1–28.
- Takahashi, S. and Yamamori, T. (2002). The pure Nash equilibrium property and the quasi-cyclic condition. *Economics Bulletin* 3, 1–6.
- Voorneveld, M. (2000). Best-response potential games. *Economics Letters* 66, 289–295.
- Young, H. P. (1993). The evolution of conventions. *Econometrica* 61, 57–84.