On the Structure of Weakly Acyclic Games ⋆

Alex Fabrikant⋆⋆, Aaron D. Jaggard⋆ ⋆ ⋆, and Michael Schapira†

Computer Science Department, Princeton University, NJ, USA;
DIMACS, Rutgers University, NJ, USA;
Departments of Computer Science, Yale University and UC Berkeley, CT, USA.

Abstract. The class of *weakly acyclic games*, which includes potential games and dominance-solvable games, captures many practical application domains. Informally, a weakly acyclic game is one where natural distributed dynamics, such as better-reply dynamics, cannot enter *inescapable oscillations*. We establish a novel link between such games and the existence of pure Nash equilibria in subgames. Specifically, we show that the existence of a *unique* pure Nash equilibrium in every *subgame* implies the weak acyclicity of a game. In contrast, we show that the existence of (potentially) *multiple* pure Nash equilibria in every subgame is *insufficient* for weak acyclicity.

1 Introduction

In many domains, convergence to a pure Nash equilibrium is a fundamental problem. In many engineered agent-driven systems that fare best when steady at pure Nash equilibrium, convergence to one is expected [7,9] to happen via *better-reply (best-reply) dynamics*: Start at some strategy profile. Players take turns, in some arbitrary order, with each player making a better reply (best reply) to the strategies of the other players, i.e., choosing a strategy that increases (maximizes) their utility, given the current strategies of the other players. Repeat this process until no player wants to switch to a different strategy, at which point we reach a pure Nash.

For better-reply dynamics to converge to a pure Nash equilibrium regardless of the initial strategy profile, a *necessary* condition is that, from every strategy profile, there exist *some* better-reply improvement path (that is, a sequence of players' better-replies) leading from that strategy profile to a pure Nash equilibrium. Games for which this property holds are called "weakly acyclic games" [10,16][1]. Both potential games [12,15] and dominance-solvable games [13]

⋆ Partially supported by the DIMACS Special Focus on Communication Security and Information Privacy.

⋆⋆ Supported by a Cisco URP grant and a Princeton University postdoctoral fellowship. alexf@cal.berkeley.edu

⋆ ⋆ ⋆ Partially supported by NSF grants 0751674 and 0753492. adj@dimacs.rutgers.edu

† Supported by NSF grant 0331548. michael.schapira@yale.edu.

[1] In some of the economics literature, the terms "weak finite improvement path property" (weak FIP) and "weak finite best response path property" (weak FBRP) are also used, for weak acyclicity under better- and best-reply dynamics, respectively.

are special cases of weakly acyclic games. In games that are not weakly acyclic, under better/best-reply dynamics, there are starting states from where the game is guaranteed to oscillate indefinitely. Moreover, the weak acyclicity of a game implies that natural decentralized dynamics (e.g., randomized better/best-reply, and no-regret dynamics) are stochastically guaranteed to reach a pure Nash [8, 16]. Thus, weakly acyclic games capture the possibility of reaching pure Nash equilibria via simple local, globally asynchronous interactions between strategic agents, independent of the starting state. We assert this is *the* realistic notion of "convergence" in most distributed systems.

Weak acyclicity has been specifically addressed in a handful of specially-structured games: in an applied setting, BGP with backup routing [1], and in a theoretical setting, games with "strategic complementarities" [3,6] (a supermodularity condition on lattice-structured strategy sets). Marden et al. [9] formulated the cooperative-control-theoretic consensus problem as a potential game (implying that it is weakly acyclic); they also defined and investigated a time-varying version of weak acyclicity.

Weak acyclicity is naturally connected to the study of the computational properties of *sink equilibria* [2,4], minimal collections of states from which best-reply dynamics cannot escape: a game is weakly acyclic if and only if all sinks are "singletons", that is, pure Nash equilibria. Unfortunately, Mirrokni and Skopalik [11] have recently examined many typical succinct representations of large games — weighted/player-specific congestion games, valid-utility games, two-sided market games, and anonymous games — and found that reliably checking weak acyclicity is extremely computationally intractable in the worst case (PSPACE-Complete). This means, inter alia, that not only can we not hope to consistently check games in these categories for weak acyclicity, but we cannot even hope to have general short "proofs" of weak acyclicity, which, once somehow found, could be tractably checked[2].

With little hope of finding robust, effective ways to consistently check weak acyclicity, we instead set out to find *sufficient* conditions for weak acyclicity: finding useable properties that imply weak acyclicity may yield better insights into at least *some* cases where we need weak acyclicity for the application.

In this work, we focus on general normal-form games. Potential games, the better-studied subset of weakly acyclic games, are known to have the following property, which we'll refer to as *subgame stability*: not only does a pure Nash equilibrium exist in the game, but a pure Nash equilibrium exists in each of its *subgames*, i.e., in each game obtained from the original game by the removal of players' strategies. Subgame stability is a useful property in many contexts. For example, in network routing games, subgame stability corresponds to the important requirement that there be a stable routing state even in the presence

[2] These two statements rely on hypotheses from computational complexity, $P \neq PSPACE$ and $NP \neq PSPACE$, both of which have remained open problems for over 40 years. While not yet formally proven, these hypotheses are almost universally believed to be true. The former is implied by the similarly near-universally-believed $P \neq NP$

of arbitrary network malfunctions [5]. We ask the following natural question: Is the extremely strong property of subgame stability *sufficient* for weak acyclicity? First, we present the following positive result for 2-player games:

Theorem 1. *In* 2*-player games, subgame stability implies weak acyclicity.*

What about games with three or more players? We show that the uniqueness of a pure Nash equilibrium in each subgame implies weak acyclicity.

Theorem 2. *If every subgame of a game Γ has a unique pure Nash equilibrium then Γ is weakly acyclic.*

In contrast, the existence of *multiple* pure Nash equilibria in subgames can lead to violations of weak acyclicity.

Theorem 3. *There are games for which subgame stability holds that are not weakly-acyclic.*

Hence, perhaps counter-intuitively, too many stable states can potentially result in the instability of local dynamics.

2 Weakly Acyclic Games and Subgame Stability

We use standard game-theoretic notation. Let Γ be a *normal-form game* with n players $1, \ldots, n$. We denote by S_i be the *strategy space* of the i^{th} player. Let $S = S_1 \times \ldots \times S_n$, and let $S_{-i} = S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$ be the cartesian product of all strategy spaces but S_i. Each player i has a *utility function* u_i that specifies i's *payoff* in every strategy-profile of the players. For each strategy $s_i \in S_i$, and every $(n-1)$-tuple of strategies $s_{-i} \in S_{-i}$, we denote by $u_i(s_i, s_{-i})$ the utility of the strategy profile in which player i plays s_i and all other players play their strategies in s_{-i}. We will make use of the following definitions.

Definition 1 (better-reply strategies). *A strategy $s'_i \in S_i$ is a* better-reply *of player i to a strategy profile (s_i, s_{-i}) if $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$.*

Definition 2 (best-reply strategies). *A strategy $s_i \in S_i$ is a* best-reply *of player i to a strategy profile $s_{-i} \in S_{-i}$ of the other players if*

$$s_i \in argmax_{s'_i \in S_i} u_i(s'_i, s_{-i}).$$

Definition 3 (pure Nash equilibria). *A strategy profile s is a* pure Nash equilibrium *if, for every player i, s_i is a best reply of i to s_{-i}*

Definition 4 (better- and best-reply improvement paths). *A* better-reply (best-reply) improvement path *in a game Γ is a sequence of strategy profiles $s^1, \ldots, s^k$ such that for every $j \in [k-1]$ (1) s^j and s^{j+1} only differ in the strategy of a single player i and (2) i's strategy in s^{j+1} is a better-reply to s^j_{-i} (best-reply to s^j_{-i} and $u_i(s^{j+1}_i, s^j_{-i}) > u_i(s^j_i, s^j_{-i})$). The* better-response dynamics (best-response dynamics) graph *for Γ is the graph on the strategy profiles in Γ whose edges are the better-reply (best-reply) improvement paths of length 1.*

We will use $\Delta R_{\Gamma}(s)$ and $BR_{\Gamma}(s)$ to denote the set of all states reachable by, respectively, better- and best-replies when starting from s in Γ.

We are now ready to define weakly-acyclic games [16]. Informally, a game is weakly acyclic if a pure Nash equilibrium can be reached from any initial strategy profile via a better-reply improvement path.

Definition 5 (weakly acyclic games). *A game Γ is* weakly acyclic *if, from every strategy profile s, there is a better-reply improvement path $s^1 \ldots, s^k$ such that $s^1 = s$, and s^k is a pure Nash equilibrium in Γ. (I.e., for each s, there's a pure Nash in $\Delta R_{\Gamma}(s)$.)*

We also coin a parallel definition based on best-reply dynamics.

Definition 6 (weakly acyclic under best-reply). *A game Γ is* weakly acyclic under best-reply *if, from every strategy profile s, there is a best-reply improvement path $s^1 \ldots, s^k$ such that $s^1 = s$ and s^k is a pure Nash equilibrium in Γ. (I.e., for each s, there's a pure Nash in $BR_{\Gamma}(s)$.)*

Weak acyclicity of either kind is equivalent to requiring that, under the respective dynamics, the game has no "non-trivial" sink equilibria [2, 4], i.e., sink equilibria containing more than one strategy profile. Conventionally, sink equilibria are defined with respect to best-response dynamics, but the original definition by Goemans et al. [4] takes into account better-response dynamics as well.

The following follows easily from definitions:

Claim. If a game is weakly acyclic under best-reply then it is weakly acyclic.

Proof. If Γ is weakly acyclic under best reply, the paths to equilibrium from each edge will still be there if we augment the state space with additional better-reply transitions. On the other hand, the game in Figure 2, mentioned, e.g., in [8], is weakly acyclic, but not weakly acyclic under best-reply.

	H	T	X
H	2,0	0,2	0,0
T	0,2	2,0	0,0
X	0,0	1,0	3,3

Fig. 1. Matching pennies with a "better-reply" escape route, but a persistent cycle under best-reply.

Curiously, all of our results apply both to weak acyclicity in its conventional better-reply sense and to weak acyclicity under best reply. Thus, unlike weak acyclicity itself, the conditions presented in this paper are "agnostic" to the better-/best-reply distinction (like the notion of pure Nash equilibria itself).

We now present the notion of subgame stability.

Definition 7 (subgames). *A* subgame *of a game Γ is a game Γ' obtained from Γ via the removal of players' strategies.*

Definition 8 (subgame stability). Subgame stability *is said to hold for a game Γ if every subgame of Γ has a pure Nash equilibrium.*

Definition 9 (unique subgame stability). Unique subgame stability *is said to hold for a game Γ if every subgame of Γ has a unique pure Nash equilibrium.*

We will also need to refer to games in which no player has two or more equally good responses to any fixed set of strategies played by the other players. Following [14], we define *strict games* as follows.

Definition 10 (strict game). *A game Γ is* strict *if, for any two distinct strategy profiles $s = (s_1, \ldots, s_n)$ and $s' = (s'_1, \ldots, s'_n)$ such that there is some $j \in [n]$ for which $s' = (s'_j, s_{-j})$ (i.e., s and s' differ only in j's strategy), then $u_j(s) \neq u_j(s')$.*

It is easy to connect unique subgame stability and strictness; the following definition is useful in proving this connection and will play a role in our main proofs as well.

Definition 11 (subgame spanned by profiles). *Given a game Γ with n players and strategy profiles $s^1, \ldots, s^k$ in Γ, the* subgame spanned by $s^1, \ldots, s^k$ *is the subgame Γ' of Γ in which the strategy space for player i is $S'_i = \{s^j_i | 1 \leq j \leq k\}$.*

Claim. Unique subgame stability implies strictness.

Proof. If a game is not strict, there are $s_j, s'_j \in S_j$ and s_{-j} such that $u_j(s_j, s_{-j}) = u_j(s'_j, s_{-j})$. Both strategy profiles in the subgame spanned by (s_j, s_{-j}) and (s'_j, s_{-j}) are pure Nash equilibria, violating unique subgame stability.

3 Sufficient Conditions for Weak Acyclicity

3.1 Subgame Stability Implies Weak Acyclicity in Two-Player Games

Theorem 1. *Every 2-player game $\Gamma = (S_1, S_2, u_{i\in\{1,2\}} : S_i \to \mathbb{R})$ such that every subgame $\Gamma' = (S'_1 \subseteq S_1, S'_2 \subseteq S_2, u_i)$ contains a pure Nash equilibrium is weakly acyclic under best-reply.*

Proof. Suppose the best-response dynamics graph has a sink component C containing multiple strategy profiles $s^1, \ldots, s^k$. Consider the subgame Γ' spanned by these strategy profiles. Let $\hat{s} = (\hat{s}_1, \hat{s}_2)$ be a pure Nash equilibrium in Γ'. By construction, there's a j such that $s^j_2 = \hat{s}_2$. Since $s^j = (s^j_1, \hat{s}_2)$ is in a non-trivial strongly connected component C, it must have at least one incoming edge and one outgoing edge connecting it with other nodes in C in the state space. These two edges must involve different players changing strategies: the same player can't play a best response and then immediately have a different best response to make. Without loss of generality, assume it has an inbound edge by player 1, from $s^{j-1} = (s_1^{j-1}, \hat{s}_2)$, and an outbound edge by player 2, to $s^{j+1} = (s_1^j, s_2^{j+1})$ (otherwise, just set $j = j+1$, and you'll still have $s_2^j = \hat{s}_2$). But $\hat{s}_1$ has to be a best response by player 1 to $\hat{s}_2$ in $\Gamma' \ni s^j$, and since $(s_1^{j-1}, \hat{s}_2)$ to $(s_1^j, \hat{s}_2)$ is a best-response in Γ (and hence in Γ'), we must have $u_1(s^j) = u_1(\hat{s})$. Thus, even in the full game Γ, where $s^{j-1} \to s^j$ is a best-reply transition, $s^{j-1} \to \hat{s}$ is also a best-reply transition. But $\hat{s}$ cannot be in C, since the best-reply links constituting C must still be present in Γ', assuring that no pure Nash of Γ' is in C, which contradicts C being a sink of Γ, with no outbound best-response edges. □

3.2 Unique Subgame Stability Implies Weak Acyclicity in All Games

Of course, Theorem 1 says nothing about n-player games for $n \geq 3$; can we guarantee weak acyclicity in these games? It turns out that we can, if every subgame has a *unique* pure Nash equilibrium.

Theorem 2. *Every game Γ that has a* unique *pure Nash equilibrium in every subgame $\Gamma' \subseteq \Gamma$ is weakly acyclic under best-reply (as are all of its subgames).*

We'll need the following technical lemma:

Lemma 1. *If s is a strategy profile in Γ, and Γ' is the subgame of Γ spanned by $BR_\Gamma(s)$, then any best-response path $s, s^1, \ldots, s^k$ in Γ' that starts at s is also a best-response path in Γ.*

Proof. Induction on the length of the path. The base case is tautological. Inductively, assume $s, \ldots, s_i$ is a best-reply path in Γ. The strategy s^{i+1} is a best-reply to s^i in Γ' by some player j. This guarantees that s^i is not a best reply by j to s^i_{-j} in Γ', let alone in Γ, so $\Gamma' \supseteq BR_\Gamma(s) \supseteq BR_\Gamma(s^i)$ must contain a best-response $\hat{s}^i_j$ to s^i_{-j} in Γ, and since s_j^{i+1} is a best-response in Γ', we are guaranteed that $u_j(\hat{s}^i_j, s^i_{-j}) = u_j(s^{i+1})$, so s^{i+1} must be a best-response in Γ.

Proof. To prove Theorem 2, assume that Γ is a game satisfying the hypotheses of the theorem, and for a subgame $\Delta \subseteq \Gamma$, denote by s_Δ the unique pure Nash equilibrium in Δ. We will proceed by induction up the semilattice of subgames of Γ. The base cases are trivial: any $1 \times \cdots \times 1$ subgame is weakly acyclic for lack of any transitions. Suppose that for some subgame Γ' of game Γ we know that every strict subgame $\Gamma'' \subsetneq \Gamma'$ is weakly acyclic.

Suppose that Γ' is not weakly acyclic: it has a state s from which its unique pure Nash $s_{\Gamma'}$ cannot be reached by best-replies. Let Γ'' be the game spanned by $BR(s)$. Consider separately the cases of (i) $s_{\Gamma'} \in \Gamma''$ and (ii) $s_{\Gamma'} \notin \Gamma''$:

Case (i): $s_{\Gamma'} \in \Gamma''$. This requires that, for an arbitrary player j with more than 1 strategy in Γ', there be a best-response path from s to some profile $\hat{s}$

where j plays the same strategy as it does in $s_{\Gamma''}$. Take one such j, and let Γ^j be the subgame of Γ' where j is restricted to playing $\hat{s}_j$ only. Since $s_{\Gamma'}$ is in Γ^j, the inductive hypothesis guarantees a best-response path in Γ^j from $\hat{s}$ to $s_{\Gamma'}$. By construction, that path must only involve best-replies by players other than j, who have the same strategy options in Γ^j as they did in Γ', so that path is also a best-reply path in Γ', assuring a best-reply path in Γ' from s to $s_{\Gamma'}$ via $\hat{s}$.

Case (ii): $s_{\Gamma'} \notin \Gamma''$. Then, Γ'''s unique pure equilibrium $s_{\Gamma''}$ must be distinct from $s_{\Gamma'}$. Since $s_{\Gamma'}$ is the only pure equilibrium in Γ', $s_{\Gamma''}$ must have an outgoing best-reply edge to some profile $\hat{s}$ in Γ'. But the inductive hypothesis ensures that $s_{\Gamma''} \in BR_{\Gamma''}(s)$, and, by the lemma, $s_{\Gamma''} \in BR_{\Gamma'}(s)$, which then ensures that $\hat{s}$ must also be in $BR_{\Gamma'}(s)$, and hence in Γ'', so $s_{\Gamma''}$ was not an equilibrium in Γ''. $\square$

4 Multiple Stable States Can Lead to Instability

What happens if we just require subgame stability, as in Theorem 1, and allow more than 2 players, as in Theorem 2? At first glance, introducing extra equilibria might seem like it would make it harder to get "stuck" in a non-trivial component of the state space with no "escape path" to an equilibrium.

This is not the case: allowing extra pure Nash equilibria in subgames actually enables the existence of non-trivial sinks.

Theorem 3. *In an n-player game, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity. This holds even for 3-player strict games.*

Proof. We will prove that the game Γ shown in Figure 4 is a counter-example that satisfies the theorem's conditions.

	c_0			c_1			c_2		
	b_0	b_1	b_2	b_0	b_1	b_2	b_0	b_1	b_2
a_0	$0,0,0$	$4,4,3$	$4,3,4$	$3,4,4$	$0,1,1$	$0,2,1$	$4,4,3$	$4,3,4$	$0,2,2$
a_1	$4,3,4$	$1,1,0$	$3,4,4$	$1,0,1$	$5,5,5$	$1,2,1$	$1,0,2$	$1,1,2$	$1,2,2$
a_2	$3,4,4$	$2,1,0$	$2,2,0$	$4,4,3$	$2,1,1$	$2,2,1$	$2,0,2$	$2,1,2$	$2,2,2$

Fig. 2. A 3-player counter-example: pure Nash equilibria exist in every subgame, but there's a persistent cycle, even under better-reply dynamics

This is a 3-player, $3 \times 3 \times 3$ game. There is a pure Nash equilibrium in the full game, $s^* = (a_1, b_1, c_1)$, with utility 5 for each of the players. There is a cycle C, every profile in which differs from s^* in at least 2 players' strategies. Any profile (a_i, b_j, c_k) that's neither s^* nor in C yields utilities (i, j, k). With utilities in C always in $3,4$, there is never an incentive for anyone to unilaterally leave the cycle C, forming a "sheath" of low-utility states separating C from the rest of the game, particularly s^*. Thus C is a persistent cycle. By construction, the game is strict and at each state in C there is a unique player who has a better reply to the current state.

Consider any subgame Γ' of this game. If Γ' contains s^*, s^* is a pure Nash equilibrium of Γ' as well.

Suppose Γ' is not the full game. In the course of cycling through C, each strategy of each player is used at least once. Thus, Γ' cannot contain all of C. If it has at least some states of C, pick one state that is in Γ', and follow the edges of C until you get to a state whose sole outbound better-response move has been "broken" by the better-response strategy being removed in Γ'. This process will terminate since C is a simple cycle in Γ that had at least one node missing in Γ'. The sole player that had an incentive to move in that state in Γ now no longer has that option, and if he has any other strategy, the resulting state cannot be in C, since C never uses more than 2 strategies of any player i in combination with any fixed s_{-i}. Thus, any other strategy is not an improvement for that player, either, and this new state is thus a pure Nash equilibrium in Γ'.

Lastly, if Γ' contains neither s^* nor any nodes of C, taking the highest-index strategy for each player yields a profile that has to be a pure Nash equilibrium, since the utilities of non-C, non-s^* profiles are just (i, j, k).

Thus, every subgame is guaranteed to have a pure Nash equilibrium, and, due to C, the game is not weakly acyclic. □

With 4 or more players, a more mechanistic approach produces counter-examples even with just 2 strategies per player:

Theorem 4. *In an n-player game for an arbitrary $n \geq 4$, the existence of pure Nash equilibria in every subgame is insufficient to guarantee weak acyclicity, even with only 2 strategies per player.*

Proof. For strategy profiles in $\{0,1\}^n$, set the utilities as follows, with all indices being modulo n:

$$\boldsymbol{u}(\boldsymbol{s}) = \begin{cases} (4, \ldots, 4) & \text{at } \boldsymbol{s} = (1, \ldots, 1) \\ (3, \ldots, 3, \underset{i\text{'th}}{2}, 3, \ldots, 3) & \text{when } s_{i-1} = s_i = 1,\ s_{-(i-1,i)} = 0 \\ (3, \ldots, 3, \underset{i+1\text{'th}}{2}, 3, \ldots, 3) & \text{when } s_i = 1,\ s_{-i} = 0 \\ \boldsymbol{s} & \text{else (for the "sheath").} \end{cases} \tag{1}$$

Similarly to Theorem 3, this plants a global pure Nash equilibrium at $(1, \ldots, 1)$, and creates a "fragile" better-response cycle. Here, the cycle alternates between profiles with edit distance $n-1$ and $n-2$ from the global pure Nash equilibrium. At every point of the cycle, the only non-sheath profiles 1 step away are its predecessor and successor on the cycle, so the cycle is persistent. Since each profile with edit distance $n-1$ from the equilibrium is covered, removing any player's 1 strategy breaks the cycle, thus guaranteeing a pure Nash equilibrium in every subgame by the same reasoning as above. □

The 4-player case of Theorem 4 is also sufficient to establish it for any n.

We note that in the 3-player, $2\times2\times2$ case, Theorem 2 still holds, even under weaker conditions:

Claim. Any strict $2\times2\times2$ game with a pure Nash equilibrium is weakly acyclic.

Proof. The pure Nash equilibrium cannot be part of the non-trivial sink, and neither can a strategy profile that differs in only one player's action from it. That leaves 4 other strategy profiles, with the possible better-response transitions forming a star in the underlying undirected graph. Since better-response links are antisymmetric ($s \to s'$ and $s' \to s$ cannot both be better-response moves), there can be no cycle among those 4 profiles, and thus no non-trivial sink components. □

5 Concluding Remarks

The connection between weak acyclicity and unique subgame stability that we present is surprising, but not immediately practicable: in most succinct game representations, there is no reason to believe that checking unique subgame stability will be tractable. In a complexity-theoretic sense, unique subgame stability is *closer* to tractability than weak acyclicity: Any reasonable game representation will have some "reasonable" representation of subgames, i.e. one in which *checking* whether a state is a pure Nash is tractable, which puts unique subgame stability in a substantially easier complexity class, $\Pi_3 P$ than the class PSPACE for which weak acyclicity is complete in many games.

We leave open the important question of finding efficient algorithms for checking unique subgame stability, which may well be feasible in particular classes of games. Also open and relevant, of course, is the question of more broadly applicable and tractable conditions for weak acyclicity.

References

1. R. Engelberg and M. Schapira. Backup routing, configuration errors and weakly acyclic games. 2009. Submitted.
2. A. Fabrikant and C. H. Papadimitriou. The complexity of game dynamics: BGP oscillations, sink equilibria, and beyond. In *Proceedings of the 19th annual ACM-SIAM symposium on Discrete algorithms*, pages 844–853, 2008.
3. J. W. Friedman and C. Mezzetti. Learning in games by random sampling. *Journal of Economic Theory*, 98:55–84, 2001.
4. M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of IEEE FOCS 2005*, pages 142–151.
5. T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Trans. Netw.*, 10(2):232–243, 2002.
6. N. S. Kukushkin, S. Takahashi, and T. Yamamori. Improvement dynamics in games with strategic complementarities. *International Journal of Game Theory*, 33(2):229–238, June 2005.

7. H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 57–66, 2008.
8. J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. Payoff-based dynamics in multi-player weakly acyclic games. *SIAM J. on Control and Optimization*, 48:373–396, 2009.
9. J.R. Marden, G. Arslan, and J.S. Shamma. Connections between cooperative control and potential games illustrated on the consensus problem. In *Proceedings of the European Control Conference*, July 2007.
10. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13:111–124, 1996.
11. V. S. Mirrokni and A. Skopalik. On the complexity of nash dynamics and sink equilibria. In *ACM Conference on Electronic Commerce*, pages 1–10, 2009.
12. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
13. H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47(6):1337–1351, 1979.
14. N. Nisan, M. Schapira, and A. Zohar. Asynchronous best-reply dynamics. In *Proceedings of the Workshop on Internet Economics*, pages 531–538, 2008.
15. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
16. H. P. Young. The evolution of conventions. *Econometrica*, 61(1):57–84, 1993.