The Negotiation Function: Subgame-perfect Equilibria in Graph Games as Fixed Points

Léonard Brice
IST Austria
Klosterneuburg, Austria
leonard.brice@ist.ac.at

Jean-François Raskin
Université libre de Bruxelles
Brussels, Belgium
jean-francois.raskin@ulb.be

Marie van den Bogaard
Université Gustave Eiffel
Champs-sur-Marne, France
marie.van-den-bogaard@univ-eiffel.fr

Nash equilibrium (NE) is a fundamental solution concept in game theory, widely used to model rational behavior among players. It is one of the most important and extensively studied concepts in the field. A strategy profile constitutes an NE if no player can deviate profitably, given the strategies of the others. Thus, NEs represent stable situations.

Unfortunately, in sequential games, NEs suffer from the problem of *non-credible threats*—see, e.g., [12, Chapter 5.5]. Specifically, in sequential games, some NEs rely on certain players *not* acting rationally following a deviation by another player; instead, they employ non-credible threats to maintain the NE and discourage deviations. The presence of non-credible threats poses a significant issue, which is why, in sequential games, the stronger concept of *subgame-perfect equilibrium* (SPE) is often preferred. A profile of strategies forms an SPE if it constitutes an NE in every subgame of the sequential game, i.e., after any possible history, and thus after any deviation, thereby eliminating non-credible threats. Thus, SPEs ensure rationality not only in players' actions but also in their planned responses to alternative scenarios.

In this talk, we discuss the contributions brought by the *negotiation function*, a powerful tool that emerged for the algorithmic study of SPEs in graph games. It was defined first in [2] (and its journal version [1]), and has been used later in [3], [4] (a journal paper gathering the results of those two papers has been recently accepted for publication in JACM), and [5]. We present the concept and its known applications, and discuss how those could be widen.

The negotiation function Let $\mathcal{G}$ be a turn-based game with infinite horizon played on a graph, with some finite number of players. Let V be the vertex set. A *requirement* is a labeling $\lambda : V \to \mathbb{R} \cup \{\pm\infty\}$. A play π (i.e., an infinite path in the underlying graph) is λ*-consistent* if each suffix $\pi_{\leq k}$ of π is a play in which the player controlling the vertex π_k gets payoff at least $\lambda(\pi_k)$. A strategy profile $\bar{\sigma}_{-i}$ (for all the players except one) is λ*-rational* if after every history h compatible with $\bar{\sigma}_{-i}$, there is a play π compatible with $\bar{\sigma}_{-i\restriction h}$ (the strategy profile $\bar{\sigma}_{-i}$ after the history h) that is λ-consistent.

The *negotiation function* is a function that takes a requirement λ, and transforms it into the (stronger or equal) requirement $\mathsf{nego}(\lambda)$, which maps every vertex v to the best payoff that the player controlling v can force, from v, against every λ-rational strategy profile.

Let $\lambda_0 : v \to -\infty$. All plays are λ_0-consistent: therefore, playing rationally with regards to the requirement λ_0, and the requirement $\lambda_1 = \mathsf{nego}(\lambda_0)$ labels every vertex with what is known in the literature as its *adversarial value*. A possible rephrasing of a classical folklore result is the following one: in a

Submitted to:
FICS 2026

large class of games, the requirement λ_1 characterises Nash equilibria, in the sense that given a play π, there exists a Nash equilibrium that generates the play π if and only if the play π is λ_1-consistent. It is shown in [2] that a similar results holds for SPEs: let $\lambda^\star$ be the least fixed point of the negotiation function (always guaranteed to exist, by Tarski's fixed point theorem). Then, there exists an SPE that generates the play π if and only if the play π is $\lambda^\star$-consistent.

Note that for $\lambda \leq \lambda'$, all λ'-consistent plays are also λ-consistent. Thus, this characterisation can also be rephrased as follows: there exists an SPE that generates the play π if and only if the play π is λ-consistent, for *some* fixed point λ of the negotiation function.

The constrained existence problem The negotiation function has emerged as an algorithmic tool to solve the SPE *constrained existence problem*: Given an initialised game $\mathcal{G}_{\upharpoonright v_0}$ and, for each player i, a lower threshold $x_i \in \mathbb{Q}$ and an upper threshold $y_i \in \mathbb{Q}$, does there exists an SPE in $\mathcal{G}_{\upharpoonright v_0}$ that generates, for each player, a payoff lying between the specified thresholds? In the literature, this problem was known to be NP-hard and EXPTIME-easy in games with parity objectives [13], and left open in games with mean-payoff objectives—two natural and classical classes of objectives. The negotiation function was first introduced to prove the decidability of the constrained existence problem for mean-payoff objectives [2, 1]. Later on, it was used to close the complexity gap for parity objectives, proving that the constrained existence problem for that class is NP-complete [3]. Then, a similar algorithm, with additional techniques, was used to prove NP-completeness of the same problem for mean-payoff objectives [4].

Interestingly, none of the algorithms presented in those two last papers uses the classical technique to compute the least fixed point of a function: namely, computing the iterations of the function from a bottom element. An algorithm of that type is suggested in a general framework in [11], and used to show tight complexity bounds in quantitative reachability games in [6] (in both cases, the negotiation function is implicitly present in the algorithm but not explicitly conceptualised); but in parity games, such an algorithm is efficient only if the number of players and of colours is small [3], and in mean-payoff games, it is not even guaranteed to terminate [1]. Instead, the non-deterministic algorithms guess a requirement and, simultaneously, a certificate of the fact that it is a fixed point of the negotiation function—by the remark made above, guessing some fixed point is sufficient, and we do not need to compute the least of them.

In games with mean-payoff objectives, it is worth noting that the negotiation function also captures ε-SPEs, a quantitative relaxation of SPEs (which are characterised by the requirement that are a fixed point of the negotiation up to ε), and that the constrained existence problem of ε-SPEs is also NP-complete [1]. This characterisation does also enable to study a variant of this problem: the *achaotic constrained existence problem*, which is motivated by rational verification [5]: Given a game $\mathcal{G}_{\upharpoonright v_0}$, and player i and a threshold t, is it true that every $\varepsilon_{\min}$-SPE in $\mathcal{G}_{\upharpoonright v_0}$ is such that player i gets at least the payoff t, where $\varepsilon_{\min}$ is the least $\varepsilon \geq 0$ such that there exists an ε-SPE in $\mathcal{G}_{\upharpoonright v_0}$? This problem is specifically relevant in games with mean-payoff objectives, where SPEs are not guaranteed to exist [15]; and in those games, using the negotiation function, it has been shown that the quantity $\varepsilon_{\min}$ exists, that it can be expressed with polynomially many bits, and that the achaotic constrained existence problem is P^{NP}-complete [5].

An unpublished extension: weak subgame-perfect equilibria An unpublished extension of the work described above is the algorithmic study of *weak subgame-perfect equilibria* (*weak SPE*). A strategy profile is a weak SPE if it constitues a *weak NE* in every subgame, that is, if in every subgame, no player can increase their payoff with a *finite* deviation of their strategy (i.e., by replacing the strategy σ_i by σ'_i such that the set $\{h \mid \sigma_i \neq \sigma'_i\}$ is finite). It is well known [7] that the definition of weak SPEs coincides to that of *very weak SPEs*, in which no player, in any subgame, has a profitable *one-shot* deviation (the set $\{h \mid \sigma_i \neq \sigma_i'\}$ is a singleton).

Weak SPEs can be captured by the *weak negotiation function*, a (simpler) variant of the negotiation function. The weak negotiation function transforms a requirement λ into the requirement $\mathsf{wnego}(\lambda)$, which maps every vertex v to the best payoff that the player controlling v can obtain, if they choose one single edge vw from v, and then an adversarial player chooses the worst possible λ-consistent play from the vertex w:

$$\mathsf{wnego}(\lambda)(v) = \max_{\text{edge } vw} \inf_{\pi\ \lambda\text{-consistent from } w} \mu_i(\pi),$$

where i is the player controlling v, and μ_i their payoff function.

We then have the following unpublished theorem.

Theorem 1 (unpublished)**.** *Let $\mathcal{G}$ be a game with parity or with mean-payoff objectives, and let π be a play in $\mathcal{G}$. Then, the play π is generated by a weak SPE if and only if there exists a fixed point λ of the weak negotiation function such that π is λ-consistent.*

This characterisation can be used to show complexity results similar as those described above for SPEs. In parity games, the following result is already known.

Theorem 2 ([8])**.** *In parity games, the constrained existence problem of weak SPEs is* NP*-complete.*

However, in mean-payoff games, a proof similar to that of [4] proves the following theorem.

Theorem 3 (unpublished)**.** *In mean-payoff games, the constrained existence problem of weak SPEs is* NP*-complete.*

Another question for which the negotiation function can be used is that of existence of equilibria. In parity games, SPEs are always guaranteed to exist [13]. Since SPEs are weak SPEs, the latter are, *a fortiori*, guaranteed to exist. Mean-payoff games, on the other hand, are not guaranteed to contain SPEs; but it is known that they always contain weak SPEs.

Theorem 4 ([9])**.** *Every mean-payoff game contains a weak SPE.*

However, a significantly shorter proof can be done using the notion of weak negotiation function: by defining the iterations $\lambda_0 : v \mapsto -\infty$, $\lambda_1 = \mathsf{wnego}(\lambda_0)$, $\lambda_2 = \mathsf{wnego}(\lambda_1)$, ..., one can show by induction that all those requirements λ_k are such that there always exists a λ_k-consistent play from every vertex, and deduce that they converge to a fixed point that also has that property, from which one can build a weak SPE.

Conjecture: randomised ε-SPEs in mean-payoff games All the works mentioned above consider a framework in which strategies are deterministic. If, now, we consider that players are allowed to randomise their strategies, the problems are significantly changed. In this new framework, it is known that the constrained existence problem of SPEs, even with reachability objectives, is undecidable [14].

As already mentioned, when players are not allowed to randomise, it is known that there exist mean-payoff games that do not contain any SPE. This result still holds when randomisation is allowed. However, in that new framework, we make the following conjecture.

Conjecture 1. *For every $\varepsilon > 0$, every mean-payoff games contains a (randomised) ε-SPE.*

We believe that a proof based on a variant of the negotiation function, similar to our proof of theorem 4, may be a possible direction, and we would like to propose this open question to the community.

References

[1] Léonard Brice, Marie van den Bogaard & Jean-François Raskin (2023): *Subgame-perfect Equilibria in Mean-payoff Games (journal version)*. *Log. Methods Comput. Sci.* 19(4), doi:10.46298/LMCS-19(4:6)2023. Available at `https://doi.org/10.46298/lmcs-19(4:6)2023`.

[2] Léonard Brice, Jean-François Raskin & Marie van den Bogaard (2021): *Subgame-Perfect Equilibria in Mean-Payoff Games*. In Serge Haddad & Daniele Varacca, editors: *32nd International Conference on Concurrency Theory (CONCUR 2021)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 203, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 8:1–8:17, doi:10.4230/LIPIcs.CONCUR.2021.8. Available at `https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CONCUR.2021.8`.

[3] Léonard Brice, Jean-François Raskin & Marie van den Bogaard (2022): *On the Complexity of SPEs in Parity Games*. In Florin Manea & Alex Simpson, editors: *30th EACSL Annual Conference on Computer Science Logic (CSL 2022)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 216, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 10:1–10:17, doi:10.4230/LIPIcs.CSL.2022.10. Available at `https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.CSL.2022.10`.

[4] Léonard Brice, Jean-François Raskin & Marie van den Bogaard (2022): *The Complexity of SPEs in Mean-Payoff Games*. In Mikolaj Bojanczyk, Emanuela Merelli & David P. Woodruff, editors: *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, *LIPIcs* 229, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 116:1–116:20, doi:10.4230/LIPICS.ICALP.2022.116. Available at `https://doi.org/10.4230/LIPIcs.ICALP.2022.116`.

[5] Léonard Brice, Jean-François Raskin & Marie van den Bogaard (2023): *Rational Verification for Nash and Subgame-Perfect Equilibria in Graph Games*. In Jérôme Leroux, Sylvain Lombardy & David Peleg, editors: *48th International Symposium on Mathematical Foundations of Computer Science, MFCS 2023, August 28 to September 1, 2023, Bordeaux, France*, *LIPIcs* 272, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 26:1–26:15, doi:10.4230/LIPICS.MFCS.2023.26. Available at `https://doi.org/10.4230/LIPIcs.MFCS.2023.26`.

[6] Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin & Marie van den Bogaard (2019): *The Complexity of Subgame Perfect Equilibria in Quantitative Reachability Games*. In: *CONCUR*, *LIPIcs* 140, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 13:1–13:16.

[7] Thomas Brihaye, Véronique Bruyère, Noémie Meunier & Jean-François Raskin (2015): *Weak Subgame Perfect Equilibria and their Application to Quantitative Reachability*. In: *24th EACSL Annual Conference on Computer Science Logic, CSL 2015, September 7-10, 2015, Berlin, Germany*, *LIPIcs* 41, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 504–518.

[8] Thomas Brihaye, Véronique Bruyère, Aline Goeminne & Jean-François Raskin (2021): *Constrained existence problem for weak subgame perfect equilibria with ω-regular Boolean objectives*. *Information and Computation* 278, p. 104594, doi:https://doi.org/10.1016/j.ic.2020.104594. Available at `https://www.sciencedirect.com/science/article/pii/S0890540120300821`. Special Issue on Ninth International Symposium on Games, Automata, Logics and Formal Verification (GandALF 2018).

[9] Véronique Bruyère, Stéphane Le Roux, Arno Pauly & Jean-François Raskin (2016): *On the existence of weak subgame perfect equilibria*. *CoRR* abs/1612.01402.

[10] Véronique Bruyère, Stéphane Le Roux, Arno Pauly & Jean-François Raskin (2017): *On the Existence of Weak Subgame Perfect Equilibria*. In: *Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, *Lecture Notes in Computer Science* 10203, pp. 145–161.

[11] János Flesch & Arkadi Predtetchinski (2016): *On refinements of subgame perfect \(\epsilon \) -equilibrium*. *Int. J. Game Theory* 45(3), pp. 523–542, doi:10.1007/s00182-015-0468-8. Available at `https://doi.org/10.1007/s00182-015-0468-8`.

[12] Martin J. Osborne (2004): *An introduction to game theory*. Oxford Univ. Press.

[13] Michael Ummels (2006): *Rational Behaviour and Strategy Construction in Infinite Multiplayer Games*. In: *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, 26th International Conference, Kolkata, India, December 13-15, 2006, Proceedings*, *Lecture Notes in Computer Science* 4337, Springer, pp. 212–223.

[14] Michael Ummels & Dominik Wojtczak (2011): *The Complexity of Nash Equilibria in Stochastic Multiplayer Games*. *Logical Methods in Computer Science* Volume 7, Issue 3, doi:10.2168/LMCS-7(3:20)2011. Available at `https://lmcs.episciences.org/1209`.

[15] Nicolas Vieille & Eilon Solan (2003): *Deterministic multi-player Dynkin games*. *Journal of Mathematical Economics* Vol.39,num. 8, pp. pp.911–929, doi:10.1016/S0304-4068(03)00021-1. Available at `https://hal-hec.archives-ouvertes.fr/hal-00464953`.