Quantifying Commitment in Nash Equilibria

Thomas A. Weber*

Forthcoming in: *International Game Theory Review*

Abstract

To quantify a player's commitment in a given Nash equilibrium of a finite dynamic game, we map the corresponding normal-form game to a "canonical extension," which allows each player to adjust his or her move with a certain probability. The commitment measure relates to the average over all adjustment probabilities for which the given Nash equilibrium can be implemented as a subgame-perfect equilibrium in the canonical extension.

JEL-Classification: C72.

Keywords: First-mover advantage; second-mover advantage; partial commitment; subgame perfection.

*Chair of Operations, Economics and Strategy, Management of Technology and Entrepreneurship Institute, Ecole Polytechnique Fédérale de Lausanne, ODY 3.01, Station 5, CH-1015 Lausanne, Switzerland. Phone: +41 21 693-0141. Email: thomas.weber@epfl.ch.

1 Introduction

How much commitment ability is needed to credibly deter market entry? How much adjustment flexibility is enough to guarantee a second-mover advantage in a pricing game? What lock-in is required to reach an advantageous outcome in a coordination game? In what follows, we develop a measure of commitment for Nash equilibria of finite-horizon dynamic games. Even in the absence of informational imperfections (assumed throughout), dynamic games can have many Nash equilibria, but almost always only one of them relies exclusively on credible threats and is therefore singled out as a subgame-perfect Nash equilibrium.[1] Which of the Nash equilibria is subgame-perfect depends on the structure of the dynamic game, in terms of the players' move order. Changing this structure does not change the set of Nash equilibria, only the specific one thought of as "credible." And any Nash equilibrium of a given normal-form representation can become subgame-perfect for a nontrivial set of corresponding dynamic games. The question we address here is how to quantify the credibility (or "commitment ability") expected from each player for a given Nash equilibrium of a finite game to be subgame-perfect.

To arrive at a measure of player-specific credibility inherent in a Nash equilibrium, independent of the structure of a particular dynamic game, we first note that a normal-form game presents an equivalence class of dynamic games which all have that same normal-form representation. For this equivalence class, we then ask: given a Nash equilibrium, *how much* commitment ability does each player expect to have in a "generic" subgame-perfect implementation of this equilibrium in extensive form? Because there is no manageable representation of all possible dynamic games belonging to a given normal form, we consider a smaller but fairly generic class of dynamic games, which we term "canonical extensions." The class of canonical extensions for a given normal-form game is a family of sequential-move games where players move once quasi-simultaneously and then, with positive probability, are subject to a termination move by one of the players. At the end of each turn, there is a positive probability for the game to end. The family of canonical extensions is fully parametrized by the vector of the players' individual continuation probabilities.

We show that for any Nash equilibrium of a finite game the set of canonical extensions for which a given Nash equilibrium is subgame-perfect has positive measure. This enables us to quantify a player's commitment in terms of a lack of expected flexibility, where the latter is given by the (expected) probability that a randomly selected canonical extension grants him an extra turn, conditional on a subgame-perfect implementation of the Nash-equilibrium at hand. Our commitment measure allows for baseline comparisons of the requirements various Nash equilibria impose on each player's credibility.

[1] A subgame-perfect Nash equilibrium (Selten 1965) is generically unique in the sense that only payoff indifference at a terminal node would give rise to multiplicity (see, e.g., Mas-Colell et al. 1995, Prop. 9.B.2). Since any small extension of the payoffs would break such indifference, the set of (finite) games with multiple subgame-perfect equilibria is of measure zero in the set of all (finite) games. Put differently, the probability that a randomly selected dynamic game has multiple subgame-perfect Nash equilibria is zero.

1.1 Literature

Heinrich Freiherr von Stackelberg (1934) pointed out the importance of move order in a duopoly. Gal-Or (1985) characterizes when, in such a duopoly setting, it is better to move first or second in terms of the firms' actions being either strategic substitutes or strategic complements. Henkel (2002) finds a possible advantage of partial commitment in duopoly games resulting from adjustment costs which may be self-imposed. Caruana and Einav (2008) endogenize commitment by making a switching cost dependent on a player's choice to wait. Weber (2014) considers the optimal choice of commitment in terms of a continuously varying adjustment cost. In the last two approaches, players have choices available for changing their commitment levels. By contrast, the question addressed here is about *how much* commitment ability is intrinsically available in a given Nash equilibrium.

Commitment also plays a role in bargaining games. For example, demand commitment in coalitional bargaining (Bennett and van Damme 1990; Selten 1992; Winter 1994) refers to the notion of players' sequentially announcing a reservation payoff for joining a coalition, usually in random order.[2] The players' commitment ability is thereby exogenous and lasts for one round of the coalition-formation game, with the last player of any given round having no commitment power at all and the first player enjoying arguably the largest commitment power (Montero and Vidal-Puga 2007). Path-dependent asymmetries, such as the random appointment of a coalition orchestrator (or "formateur") in each round with hold-up power also skews the ex-post payoff distribution (Breitmoser 2009).

The commitment measure we construct here can be used *ex post* for any Nash equilibrium in a finite-horizon dynamic game with endogenous commitment. We remain unconcerned with the question of how precisely the parties' commitment might arise, which commitment devices they might be using, and so forth. More specifically, we think of commitment as a lack of flexibility, whereby a player's flexibility in a given Nash equilibrium is measured as an "expected continuation probability" in a randomly selected canonical extension, conditional on the fact that this dynamic game may serve to implement the said Nash equilibrium.

1.2 Outline

In Section 2, we assume that a finite dynamic game is represented in (reduced) normal form (Fudenberg and Tirole 1991) and use the primitives of this representation to introduce a "canonical extension," which contains a flexibility parameter for each player. Section 3 shows how a pure-strategy Nash equilibrium of the normal-form game can be mapped to an equilibrium of the canonical extension for suitable values of the flexibility parameter. In Section 4, we argue that the flexibility parameter can be used to measure the players' commitment levels inherent in any given Nash equilibrium of the normal-form game. Section 5 concludes.

[2]The idea of the demand commitment game originated with Reinhard Selten (Bennett and van Damme 1990, p. 2).

2 Models

2.1 Normal-Form Game

Consider the normal-form game $\Gamma = \{\mathcal{N}, \{\mathcal{A}_i\}_{i\in\mathcal{N}}, \{u_i\}_{i\in\mathcal{N}}\}$ for $N \geq 2$ players in the player set $\mathcal{N} = \{1, \ldots, N\}$. Each player $i \in \mathcal{N}$ has a nonempty finite action set $\mathcal{A}_i$ and a real-valued payoff function $u_i : \mathcal{A} \to \mathbb{R}$, where $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$ is the set of all strategy profiles. A pure-strategy profile $a^* \in \mathcal{A}$ is a Nash equilibrium of Γ if

$$a^* \in \mathcal{R}(a^*), \tag{1}$$

where, for any strategy profile $a \in \mathcal{A}$, the (static) best-response correspondence is

$$\mathcal{R}(a) \triangleq \prod_{i\in\mathcal{N}} \mathcal{R}_i(a_{-i}) = \prod_{i\in\mathcal{N}} \arg\max_{a_i\in\mathcal{A}_i} u_i(a_i, a_{-i}), \tag{2}$$

using the standard notational convention that $a = (a_i, a_{-i})$ for any $i \in \mathcal{N}$. The definition of a mixed-strategy Nash equilibrium is more general and is obtained by replacing each action set $\mathcal{A}_i$ with the corresponding simplex $\Delta(\mathcal{A}_i)$ of probability distributions over actions. To simplify the exposition, we first assume that there exists at least one pure-strategy Nash equilibrium a^* of Γ, postponing the discussion of more general situations to Section 5.[3]

Example 1. Consider the standard battle-of-the-sexes game (Luce and Raiffa 1957, pp. 90–91; Osborne and Rubinstein 1984, pp. 15–16), where each of two players, whom we call Ann (player 1) and Bert (player 2), can choose to either go dancing ('D') or go to the movies ('M'). Whenever the players choose different actions, their payoffs vanish. Otherwise the players obtain the payoffs $u_1(D, D) = u_2(M, M) = 2$ and $u_1(M, M) = u_2(D, D) = 1$, respectively. Since each player i's static best-response correspondence is $\mathcal{R}_i(a_{-i}) \equiv \{a_{-i}\}$, this 'coordination game' has two pure-strategy Nash equilibria $a^* \in \{(D, D), (M, M)\}$; it also has a mixed-strategy Nash equilibrium where each player chooses his or her preferred action with probability $2/3$.

2.2 Canonical Extension

Let $p = (p_1, \ldots, p_N) \in \Delta_N$ be a vector of continuation probabilities in the N-dimensional simplex

$$\Delta_N \triangleq \{(p_1, \ldots, p_N) \in \mathbb{R}^N_+ : p_1 + \cdots + p_N \leq 1\}. \tag{3}$$

The *canonical extension* $\hat{\Gamma}(p)$ of the normal-form game Γ is a two-stage dynamic game which extends over the time periods $t \in \{0, 1\}$. At time $t = 0$, all players play a simultaneous-move stage game of the form Γ. At the end of that first period, the game ends with probability

$$p_0 \triangleq 1 - (p_1 + \cdots + p_N). \tag{4}$$

Otherwise, with probability p_i, player i is allowed to adjust his move in the second period, at time $t = 1$; all other players are stuck with their first-period actions. At the end of

[3]The Kakutani fixed-point theorem guarantees that Γ has a (mixed-strategy) Nash equilibrium (Nash 1950).

the second period, the game ends and all players obtain their (undiscounted) payoffs, based on the strategy profile $a = (a_i, a_{-i})$ for player i's second-period choice a_i and the profile a_{-i} of the other players' first-period choices.

The canonical extension can be viewed as a dynamic perturbation of Γ where each player's influence may spill into a second period, thus potentially limiting the commitment power for any given player. For $p = 0$, the game terminates with probability 1 at the end of the first period and we obtain the original game. That is,

$$\hat{\Gamma}(0) = \Gamma. \tag{5}$$

A subgame-perfect equilibrium of the canonical extension $\hat{\Gamma}(p)$ is obtained by backward-induction. At time $t = 1$, conditional on the first-period strategy profile $a = (a_i, a_{-i}) \in \mathcal{A}$, player i—upon being able to adjust his move—chooses an element of his best response,

$$r_i(a_{-i}) \in \mathcal{R}_i(a_{-i}) \triangleq \arg\max_{\hat{a}_i \in \mathcal{A}_i} u_i(\hat{a}_i, a_{-i}). \tag{6}$$

Hence, at time $t = 0$, each player i obtains the expected utility

$$U_i(a_i, a_{-i}; p) \triangleq p_0 u_i(a_i, a_{-i}) + p_i u_i(r_i(a_{-i}), a_{-i}) + \sum_{j \in \mathcal{N} \setminus \{i\}} p_j u_i(a_i, r_j(a_i, (a_{-i})_{-j}), (a_{-i})_{-j}),$$

where

$$(a_{-i})_{-j} \triangleq (a_l)_{l \in \mathcal{N} \setminus \{i,j\}}, \quad i, j \in \mathcal{N},\ i \neq j,$$

denotes the strategy profile of the players other than players i and j. With probability p_0 player i obtains his standard stage-game payoff, and with probability p_i he gets to observe the other players' strategy profile and choose a best response; otherwise, player i can anticipate the expected payoff impact of other players' reacting to the strategy profile they may observe upon being given the opportunity to adjust their actions. Fig. 1 illustrates the payoffs in the canonical extension $\hat{\Gamma}(p)$ for a two-player game where each player has two possible actions. After playing a simultaneous-move game, with probability p_i player i gets to move again, thus instead of the previously played a_i choosing a best response $r_i(a_{-i})$ to the other player's action a_{-i}.

Remark 1 (Stackelberg Variants). A pure Stackelberg variant of $\hat{\Gamma}(p)$ with player i moving last is obtained for $p = (1, \ldots, 1) - e_i$, where e_i is the i-th unit vector in the standard Euclidean base of $\mathbb{R}^N$. Similarly, a Stackelberg variant of $\hat{\Gamma}(p)$ where player i moves first is obtained for $p = (p_i, p_{-i})$ where $p_i = 0$ and $p_{-i} \gg 0$ is such that $\sum_{j \in \mathcal{N} \setminus \{i\}} p_j = 1$, so the probability of stopping at the end of the first period vanishes ($p_0 = 0$).[4] One of the other players will surely terminate the game. For $N = 2$ the dynamic game $\hat{\Gamma}(p)$ can therefore embed any pure player order. For $N > 2$, the dimensionality of the simplex Δ_N is not high enough for p to be able to continuously vary among all possible player orders (of which there are $N! > N$ for $N > 2$). Still, it is possible to analyze within $\hat{\Gamma}(p)$ the comparative statics when any given player's move position is varied from first to last or vice-versa.

[4] We use standard notation for vector inequalities, where for $p = (p_1, \ldots, p_N)$ and $\hat{p} = (\hat{p}_1, \ldots, \hat{p}_N)$ it is $p \geq \hat{p} \ \Leftrightarrow\ (p_i \geq \hat{p}_i$ for all $i \in \mathcal{N})$, $p > \hat{p} \ \Leftrightarrow\ (p \geq \hat{p}$ and $p_i > \hat{p}$ for at least one $i \in \mathcal{N})$, and $p \gg \hat{p} \ \Leftrightarrow$ $(p_i > \hat{p}_i$ for all $i \in \mathcal{N})$.

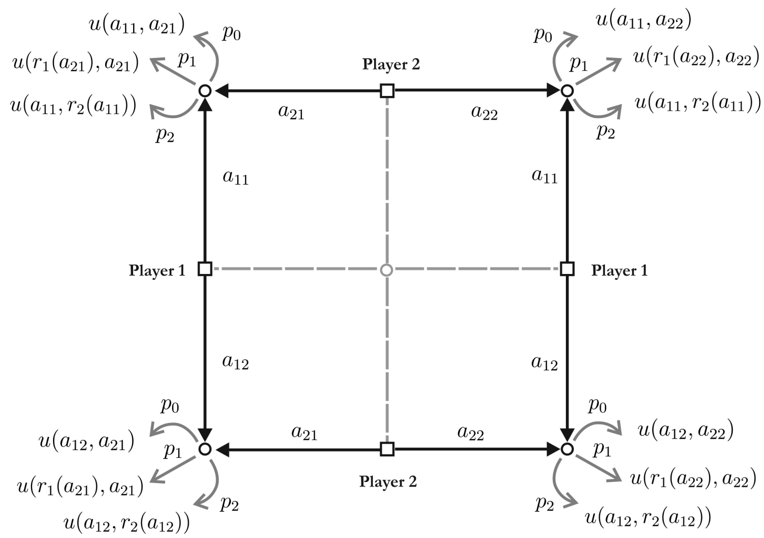

Figure 1: Extensive-form representation of the canonical extension $\hat{\Gamma}(p)$ for $N = 2$ and $\mathcal{A} = \{a_{11}, a_{12}\} \times \{a_{21}, a_{22}\}$, where $u = (u_1, u_2)$ and $p_0 = 1 - p_1 - p_2$.

Given any $p \in \Delta_N$, the canonical extension $\hat{\Gamma}(p)$ can be viewed as the original normal-form game Γ with the payoff functions $u_i(\cdot)$ replaced by $U_i(\cdot; p)$, for all $i \in \mathcal{N}$. Thus, we consider the parametrized family of *canonical normal-form games*,

$$\Gamma(p) \triangleq \{\mathcal{N}, \{\mathcal{A}_i\}_{i \in \mathcal{N}}, \{U_i(\cdot; p)\}_{i \in \mathcal{N}}\}, \quad p \in \Delta_N.$$

In particular, a strategy profile $a^* = (a_1^*, \ldots, a_N^*)$ is a Nash equilibrium of the normal-form game $\Gamma(p)$ if and only if it is a (subgame-perfect) Nash equilibrium of the canonical extension $\hat{\Gamma}(p)$ for each player $i \in \mathcal{N}$ to play a_i^* in the first period and to play $a_i^* = r_i(a_{-i}^*)$ provided he gets to move again in the second period.

Lemma 1. *For any $p \in \Delta_N$, there is a one-to-one correspondence between the Nash equilibria of the canonical normal-form game $\Gamma(p)$ and the subgame-perfect Nash equilibria of the canonical extension $\hat{\Gamma}(p)$.*

The preceding result allows us to think of the canonical extension $\hat{\Gamma}(p)$ as equivalent to the canonical normal-form game $\Gamma(p)$. Each one also presents an embedding of the original normal-form game, as $\Gamma = \Gamma(0) = \hat{\Gamma}(0)$.

Example 2. In the canonical extension $\hat{\Gamma}(p)$ of the battle-of-the-sexes game in Ex. 1, Ann and Bert effectively choose their activities in turn. As noted in Remark 1, the value of $p = (p_1, p_2) \in \Delta_2$ determines who starts and who terminates the game. For example, when $p = (1, 0)$ Ann effectively starts the game, while Bert's move concludes their interaction. Consider now the strategy profile $a^* = (D, D)$ in the canonical normal-form game $\Gamma(p)$. Ann would never want to deviate because

$$U_1(a^*; p) = 2p_0 + 2p_1 + 2p_2 \equiv 2 \geq 2p_1 + p_2 = U_1(M, D; p),$$

for all $p \in \Delta_2$. However, Bert would prefer a deviation, provided that Ann's continuation probability is large enough. Indeed,

$$U_2(M, D; p) = 0 + p_2 + 2p_1 \geq 1 = p_0 + p_2 + p_1 = U_2(a^*; p),$$

if and only if $p_1 \geq (1 - p_2)/2$; this condition is automatically satisfied if $p_1 \geq 1/2$. In other words, high flexibility for Ann destroys her commitment ability.

3 Sequential Implementation

For any given $p \in \Delta_N$, we say that a Nash equilibrium $a^* = (a_i^*)_{i \in \mathcal{N}}$ of the simultaneous-move game Γ can be *implemented sequentially* in the canonical extension $\hat{\Gamma}(p)$ if and only if a^* is a Nash equilibrium of the canonical normal-form game $\Gamma(p)$. As discussed in the last section, this Nash equilibrium then induces a subgame-perfect equilibrium in $\hat{\Gamma}(p)$.

Remark 2 (Comparison with Repeated Games). The idea of dynamically implementing static Nash equilibria is not new. Indeed, one can implement the payoffs of any pure-strategy Nash equilibrium of Γ as average payoffs of a repeated game Γ^∞ where in each period the players play the simultaneous-move stage game Γ and payoffs are discounted from period to period. In such a repeated game, a subgame-perfect strategy profile is for players to ignore the history of past actions and simply play the Nash equilibrium a^* in each period. However, this approach affords no insights into the relationship of the normal-form game Γ to its sequential-move variants.[5]

For any player $i \in \mathcal{N}$, let

$$\mathcal{P}_i(a^*) \triangleq \left\{ \hat{p} \in \Delta_N : \max\{U_i(a_i, a_{-i}^*; \hat{p}) : a_i \in \mathcal{A}_i\} \leq u_i(a^*) \right\} \tag{7}$$

be the set of all continuation probabilities p such that player i's deviation in $\Gamma(p)$ is *not* profitable relative to the equilibrium payoff at the Nash equilibrium a^* in Γ.

Proposition 1. *A Nash equilibrium a^* of Γ can be implemented sequentially in $\hat{\Gamma}(p)$ if and only if*

$$p \in \mathcal{P}(a^*), \tag{8}$$

where $\mathcal{P}(a^*) \triangleq \cap_{i \in \mathcal{N}} \mathcal{P}_i(a^*)$.

Proof. Fix $p \in \Delta_N$, and let a^* be a Nash equilibrium of Γ. In the canonical extension $\hat{\Gamma}(p)$, any given player $i \in \mathcal{N}$ prefers the Nash equilibrium action a_i^* to any other available action $a_i \in \mathcal{A}_i$ (at least weakly) if and only if

$$U_i(a_i, a_{-i}^*; p) \leq U_i(a_i^*, a_{-i}^*; p), \quad a_i \in \mathcal{A}_i. \tag{9}$$

Since $U_i(a_i^*, a_{-i}^*; p) = u_i(a^*)$, the last relation is equivalent to p being in the set $\mathcal{P}_i(a^*)$ as specified in Eq. (7). In a sequential implementation of a^*, no player can have an incentive to deviate from the equilibrium strategy profile, restricting the corresponding probability vectors p to lie in the intersection $\mathcal{P}(a^*) = \cap_{i \in \mathcal{N}} \mathcal{P}_i(a^*)$, completing our proof. □

Consider now the measure of the set of continuation probabilities,

$$\|\mathcal{P}(a^*)\| \triangleq \int_{\mathcal{P}(a^*)} dp.$$

[5] Infinitely repeated games feature generic equilibrium multiplicities due to well-known folk theorems (see, e.g., Friedman 1971; Fudenberg and Maskin 1986) that rely on a variety of credible out-of-equilibrium threats. To bypass such complications we opt here for simple sequential implementations in two periods.

The sequential implementation of a^* is called *strong* if the set of continuation probabilities has positive measure. Our main result is that any Nash equilibrium of Γ has a (strong) sequential implementation.

Corollary 1. *Any Nash equilibrium a^* of Γ has a strong sequential implementation in $\hat{\Gamma}(p)$ for $p \in \mathcal{P}(a^*)$; and, $\|\mathcal{P}(a^*)\| > 0$.*

Proof. Let a^* be a Nash equilibrium of Γ. By Eq. (5) this Nash equilibrium is trivially implementable in $\hat{\Gamma}(0)$ (i.e., $0 \in \mathcal{P}(a^*)$). Consider now the possible deviation a_i for player $i \in \mathcal{N}$, such that his action is *not* in his best response to the equilibrium strategy profile a^*_{-i}; that is, let [6]

$$a_i \in \bar{\mathcal{R}}_i(a^*_{-i}) \triangleq \mathcal{A}_i \setminus \mathcal{R}_i(a^*).$$

Thus, because $\mathcal{A}_i$ is finite, there exists an $\varepsilon_i > 0$ such that

$$\max_{a_i \in \mathcal{A}_i} U_i(a_i, a^*_{-i}; 0) + \varepsilon_i \leq U_i(a^*_i, a^*_{-i}; 0).$$

By the maximum theorem (Berge 1959) the envelope $m_i(p) \triangleq \max_{a_i \in \mathcal{A}_i} U_i(a_i, a^*_{-i}; p)$ is continuous in p, so that there exists a $\delta_i > 0$ for which:[7]

$$\|p\|_\infty \in [0, \delta_i] \quad \Rightarrow \quad |m(0) - m(p)| \leq \varepsilon_i.$$

This implies that $\mathcal{P}_i(a^*) \subset [0, \delta_i]^N$. If we set $\delta \triangleq \min_{i \in \mathcal{N}} \delta_i > 0$ and repeat the preceding arguments for all other players, one obtains that $\mathcal{P}(a^*) \subset [0, \delta]^N$, which implies that $\|\mathcal{P}(a^*)\| \geq \delta^N > 0$, completing the proof. $\square$

For any $i, j \in \mathcal{N}$ and any Nash equilibrium $a^* = (a^*_i, a^*_{-i}) \in \mathcal{A}$ of Γ, consider the mapping $q_{i,j}(\cdot; a^*) : \mathcal{A}_i \to \mathbb{R}$, with

$$q_{i,j}(a_i; a^*) \triangleq \begin{cases} \dfrac{u_i(a_i, r_j(a_i, (a^*_{-i})_{-j}), (a^*_{-i})_{-j}) - u_i(a_i, a^*_{-i})}{u_i(a^*) - u_i(a_i, a^*_{-i})}, & \text{if } a_i \in \bar{\mathcal{R}}_i(a^*_{-i}) \text{ and } i \neq j, \\ 1, & \text{otherwise,} \end{cases}$$

for all $a_i \in \mathcal{A}_i$. The quotient $q_{i,j}(a_i; a^*)$ quantifies player i's payoff variation if player j is able to adjust to i's deviation a_i (resulting in j's choosing the best response $r_j(a_i, (a^*_{-i})_{-j})$ instead of a^*_j), relative to player i's benefit of playing the Nash-equilibrium action a^*_i instead of unilaterally deviating to a_i. The underlying assumption is that all agents other than i and j are playing according to the Nash-equilibrium strategy profile $(a^*_{-i})_{-j}$. Intuitively, the quotient $q_{i,j}(a_i; a^*)$ quantifies the ratio of payoff variations between bilateral deviations (first i, then j) and unilateral deviations (for i only).

Proposition 2. *The set of continuation probabilities for which a sequential implementation of the Nash equilibrium a^* is possible can be represented in the form:*

$$\mathcal{P}(a^*) = \{p \in \Delta_N : f(p; a^*) \leq 1\}, \tag{10}$$

[6] If $\bar{\mathcal{R}}_i(a^*) = \emptyset$, then $\mathcal{R}_i(a^*) = \mathcal{A}_i$, which implies that $\mathcal{P}_i(a^*) = \Delta_N$, so that player i has no part in restricting $\mathcal{P}(a^*)$ and can therefore be neglected. If $\bar{\mathcal{R}}_i(a^*) = \emptyset$ for all $i \in \mathcal{N}$, then $\|\mathcal{P}(a^*)\| = \|\Delta_N\| > 0$.

[7] Because of the norm-equivalence in finite-dimensional Euclidean spaces, the precise choice of the norm is qualitatively unimportant. For simplicity, we here make use of the maximum-norm, defined as $\|p\|_\infty \triangleq \max\{p_1, \ldots, p_N\}$, for all $p = (p_1, \ldots, p_N) \in \Delta_N$.

where $f(p;a^*)$ is piecewise linear and convex in p with $f(0;a^*)=0$, and

$$f(p;a^*) = \max_{i\in\mathcal{N}} \left\{ \max_{a_i\in\bar{\mathcal{R}}_i(a^*_{-i})} \sum_{j\in\mathcal{N}} q_{i,j}(a_i;a^*)\, p_j \right\}, \quad p\in\Delta_N. \tag{11}$$

Proof. Given any Nash equilibrium a^* of Γ and any $p\in\Delta_N$, the implementability condition in Eq. (9) is equivalent to

$$p_0 u_i(a_i,a^*_{-i}) + p_i u_i(a^*) + \sum_{j\in\mathcal{N}\setminus\{i\}} p_j u_i(a_i, r_j(a_i,(a^*_{-i})_{-j}),(a^*_{-i})_{-j}) \le u_i(a^*), \quad a_i\in\bar{\mathcal{R}}_i(a^*_{-i}).$$

Using the fact that $p_0 = 1-\sum_{i\in\mathcal{N}} p_i$ (see Eq. (4)) and that $u_i(a^*)-u_i(a_i,a^*_{-i})>0$ for all $a_i\in\bar{\mathcal{R}}_i(a^*_{-i})$, the preceding inequality can be rewritten in the form

$$p_i + \sum_{j\in\mathcal{N}\setminus\{i\}} p_j \frac{u_i(a_i,r_j(a_i,(a^*_{-i})_{-j}),(a^*_{-i})_{-j}) - u_i(a_i,a^*_{-i})}{u_i(a^*)-u_i(a_i,a^*_{-i})} \le 1, \quad a_i\in\bar{\mathcal{R}}_i(a^*_{-i}).$$

Hence, player i's expected equilibrium payoff $u_i(a^*) = U_i(a^*;p)$ exceeds the deviation payoff $U_i(a_i,a^*_{-i};p)$ (at least weakly) if and only if

$$\sum_{j\in\mathcal{N}} q_{i,j}(a_i;a^*)\, p_j \le 1. \tag{12}$$

Maximizing the left-hand side of the last condition over all $a_i\in\bar{\mathcal{R}}_i(a^*_{-i})$ and then over all $i\in\mathcal{N}$ yields the expression for $f(p;a^*)$ and the claimed representation of $\mathcal{P}(a^*)$ in Eq. (10). Finally, as an upper envelope of linear functions, $f(\cdot;a^*)$ is naturally convex. This completes the proof. □

As long as player i's possible deviations a_i are not in his best response $\mathcal{R}_i(a^*_{-i})$ (i.e., they should lie in the complement $\bar{\mathcal{R}}_i(a^*_{-i})$), the denominator of $q_{i,j}(a_i;a^*)$ is strictly positive. However, the numerator may well be negative, indicating an additional disbenefit from player j's ability to also deviate. This in turn points to a first-mover advantage and facilitates the sequential implementation of a^* in $\hat{\Gamma}(p)$ for a particular vector of continuation probabilities p. In other words, the more negative the $q_{i,j}(a_i;a^*)$ (for $j\neq i$), the more likely the Nash equilibrium survives a perturbation from Γ to $\hat{\Gamma}(p)$.

Lottery Interpretation. If we introduce the discrete lottery

$$L_i(a_i;a^*,p) = \left[p_0, 0; [p_j, q_{i,j}(a_i;a^*)]_{j=1}^N\right],$$

with the $N+1$ payoffs $0, q_{i,1}(a_i;a^*),\dots,q_{i,N}(a_i;a^*)$ that occur with probabilities $p_0,p_1,\dots,p_N$, respectively,[8] then

$$U_i(a_i,a^*_{-i};p) \le U_i(a^*;p) \quad \Leftrightarrow \quad \mathbb{E}\left[L_i(a_i;a^*,p)\right] \le 1.$$

In other words, for the Nash equilibrium a^* to be sequentially implementable, any player i's expected relative payoff variations from bilateral reactions to unilateral deviations cannot exceed unity.

[8] We recall that $p_0 = 1-(p_1+\cdots+p_N)$ by Eq. (4), so that there is no additional dependence on p_0.

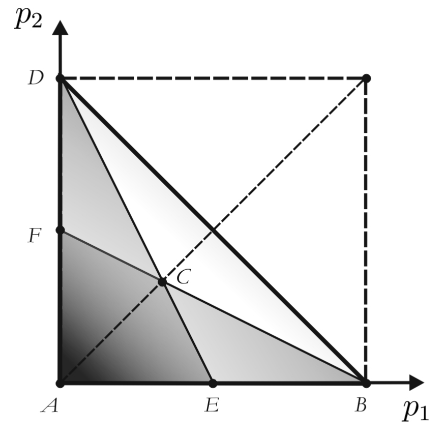

Figure 2: Regions of sequential implementability (as subsets of Δ_2) in Ex. 3.

Corollary 2. *For any $i \in \mathcal{N}$, the set $\mathcal{P}_i(a^*)$ is convex.*

Proof. Consider $p, \hat{p} \in \mathcal{P}_i(a^*)$ and $\theta \in (0,1)$. Then for any $a_i \in \bar{\mathcal{R}}_i(a^*_{-i})$, it is

$$\sum_{j \in \mathcal{N}} q_{i,j}(a_i; a^*_{-i})(\theta p_j + (1-\theta)\hat{p}_j) = \theta \sum_{j \in \mathcal{N}} q_{i,j}(a_i; a^*_{-i})\, p_j + (1-\theta) \sum_{j \in \mathcal{N}} q_{i,j}(a_i; a^*_{-i})\, \hat{p}_j \leq 0,$$

which by inequality (12) implies that

$$U_i(a^*; \theta p + (1-\theta)\hat{p}) \geq U_i(a_i, a^*_{-i}; \theta p + (1-\theta)\hat{p}), \quad a_i \in \bar{\mathcal{R}}_i(a^*_{-i}),$$

so that, by definition, $\theta p + (1-\theta)\hat{p} \in \mathcal{P}_i(a^*)$, which establishes the claim.

□

Example 3. Using Prop. 2 the Nash equilibria of the simultaneous-move battle-of-the-sexes game in Ex. 1 can be implemented sequentially. From the discussion in Ex. 2 we obtain that the Nash equilibrium $a^* = (D, D)$ of Γ is also a Nash equilibrium of the canonical extension $\hat{\Gamma}(p)$ if and only if p lies in the set $\mathcal{P}(a^*) = \mathcal{P}_1(a^*) \cap \mathcal{P}_2(a^*)$, where $\mathcal{P}_1(a^*) = \Delta_2$ and $\mathcal{P}_2(a^*) = \{(\hat{p}_1, \hat{p}_2) \in \Delta_2 : \hat{p}_1 \leq (1 - \hat{p}_2)/2\}$. Hence, $\mathcal{P}(a^*) = \mathcal{P}_2(a^*)$. By symmetry, for $\hat{a}^* = (M, M)$ we obtain that $\mathcal{P}(\hat{a}^*) = \{(\hat{p}_1, \hat{p}_2) \in \Delta_2 : \hat{p}_2 \leq (1 - \hat{p}_1)/2\}$. Thus, both pure-strategy Nash equilibria in the battle-of-the-sexes game of Ex. 1 can be implemented sequentially in $\hat{\Gamma}(p)$ as long as $p = (p_1, p_2)$ lies in the region $AECF$; see Fig. 2. More specifically, to sequentially implement Ann's preferred Nash equilibrium (D, D) it is enough if p lies in the region $AE(C)D$, while Bert's preferred equilibrium (M, M) can be sequentially implemented if p is in $AB(C)F$. By embedding the normal-form game into the dynamic framework, one learns that in order for Ann to obtain her favorite choice (D), her probability (p_1) of being able to adjust her strategy cannot exceed half of Bert's probability of not being able to adjust his choice $(1 - p_2)$. Thus, it is *not* necessary that player 2's flexibility is smaller than player 1's because for any $p_1 \in (0, 1)$, there exists a $p_2 < p_1$ that still allows for a dynamic implementation of player 1's preferred equilibrium.[9] By continuously varying p over the parameter space Δ_2 we obtain the implementable equilibria in the form of an upper semicontinuous correspondence, and

[9] Any player i's ability to commit decreases in p_i.

thus a conceptually richer picture of commitment than can be obtained by just comparing the two Stackelberg versions. The parameter space captures the underlying continuity of commitment levels (Weber 2014).

4 A Measure of Commitment

The nonempty set of continuation-probability vectors $\mathcal{P}(a^*)$ for any given Nash equilibrium a^* of the normal-form game Γ serves as the basis for a measure of the players' commitment ability.

Expected Flexibility. For each player i, the higher the continuation probability p_i that sustains a^*, the higher his *expected flexibility* in this Nash equilibrium, defined as

$$\varphi_i(a^*) \triangleq \frac{\int_{\mathcal{P}(a^*)} p_i dp}{\|\mathcal{P}(a^*)\|}, \tag{13}$$

where $\|\mathcal{P}(a^*)\| = \int_{\mathcal{P}(a^*)} dp > 0$ by Cor. 1. The flexibility measure $\varphi_i(a^*) \in [0, 1]$ represents player i's expected continuation probability in a random instance of the canonical extension $\hat{\Gamma}(p)$ conditional on being in a sequential implementation of the Nash equilibrium a^* of Γ.

Lemma 2. *For any Nash equilibrium a^* of Γ and any player $i \in \mathcal{N}$: $\varphi_i(a^*) \in [0, 1/2]$.*

Proof. Fix $i \in \mathcal{N}$. Given that $p \in \Delta_N$, player i's flexibility p_i can be largest (i.e., $p_i = 1$) if $p_0 = 0$ and all other players' flexibilities vanish (i.e., $p_j = 0$ for all $j \in \mathcal{N} \setminus \{i\}$). By Cor. 2, the set $\mathcal{P}_i(a^*)$ is convex, so that if $p_i = \hat{p}_i \in (0, 1)$ is feasible, then $p_i \in [0, \hat{p}_i)$ must also be feasible. By Cor. 1, the set $\mathcal{P}(a^*)$ has a positive measure, which implies there exists an $\varepsilon \in (0, 1/N)$, so that

$$p \in \mathcal{B}_{N,\varepsilon} \triangleq \{\hat{p} \in \Delta_N : \|\hat{p}\|_\infty \leq \varepsilon\} \quad \Rightarrow \quad p \in \mathcal{P}(a^*).$$

Since $\mathcal{B}_{N,\varepsilon} \subset \mathcal{P}(a^*)$, the largest possible value for p_i is $\bar{p}_{i,\varepsilon} = 1 - (N - 1)\varepsilon$. If we now set $\hat{p}_i$ equal to the largest value in the i-th direction of elements in the set $\mathcal{P}(a^*)$,

$$\hat{p}_i = \sup \{p_i \in [0, 1] : (p_i, p_{-i}) \in \mathcal{P}(a^*)\},$$

then $\hat{p}_i < 1$. By choosing ε such that $0 < \varepsilon < (1 - \hat{p}_i)/(N - 1)$, we obtain that $\hat{p}_i < \bar{p}_{i,\varepsilon}$, which in turn implies that

$$\varphi_i(a^*) < \frac{\int_{[0,\varepsilon]^{N-1} \times [0,\bar{p}_{i,\varepsilon}]} p_i dp}{\|[0, \varepsilon]^{N-1} \times [0, \bar{p}_{i,\varepsilon}]\|} = \frac{1 - (N - 1)\varepsilon}{2} \uparrow \frac{1}{2} \quad (\text{for } \varepsilon \downarrow 0^+),$$

concluding our proof. $\square$

The proof of La. 2 shows that 1/2 is in fact a *tight* upper bound for any player i's expected flexibility. It turns out that the sum of all players' expected flexibilities cannot exceed twice the expected flexibility of any one player.

Lemma 3. *For any Nash equilibrium a^* of Γ: $\sum_{i \in \mathcal{N}} \varphi_i(a^*) \leq 1$.*

Proof. Let a^* be a Nash equilibrium of Γ. By the definition of the players' expected flexibility $\varphi_i(a^*)$ in Eq. (13) and the definition of the set Δ_N of admissible probability vectors in Eq. (3), it is

$$\sum_{i\in\mathcal{N}} \varphi_i(a^*) = \frac{\int_{\mathcal{P}(a^*)} \left(\sum_{i\in\mathcal{N}} p_i\right) dp}{\|\mathcal{P}(a^*)\|} \leq \frac{\int_{\mathcal{P}(a^*)} dp}{\|\mathcal{P}(a^*)\|} = 1,$$

which establishes the claim. □

Commitment Measure. We are now ready to introduce player i's *measure of commitment* (*in the Nash equilibrium* a^*) as follows:

$$\kappa_i(a^*) \triangleq \max\{0, 1 - N\varphi_i(a^*)\}. \tag{14}$$

This measure of commitment takes on values between the tight bounds 0 and 1. It can be interpreted as the renormalized probability with which player i's decision cannot be revised in a random canonical extension that allows for an implementation of the Nash equilibrium a^*. The renormalization is such that player i's commitment vanishes whenever his expected flexibility exceeds $1/N$, which for $N = 2$ cannot happen at all because φ_i is, by La. 2, limited to values in $[0, 1/2]$. The reason for the normalization becomes clear with the next result and the example thereafter.

Lemma 4. *For any Nash equilibrium a^* of Γ, it is $(1/N)\sum_{i\in\mathcal{N}} \kappa_i(a^*) \in [0, 1]$.*

Proof. The claim follows immediately when applying La. 3 to Eq. (14). □

The players' average commitment is bounded tightly by 0 and 1, and therefore the measure retains full informativeness as the number of players goes up.[10]

Example 4 (**Full Sequential Implementability and Reference Commitment**). Consider a *trivial* game (e.g., with constant payoffs for each player) where any Nash equilibrium a^* can be implemented sequentially in the dynamic extension $\hat{\Gamma}(p)$ for *any* $p \in \Delta_N$. Given that $\|\Delta_N\| = 1/N!$, we find that

$$\kappa_i(a^*) = \max\{0, 1-N\varphi_i(a^*)\} = 1-\frac{N\int_{\Delta_N} p_i\, dp}{\|\Delta_N\|} = 1-\frac{N/(N+1)!}{(1/N!)} = \frac{1}{N+1} = \varphi_i(a^*) > 0,$$

for any $i \in \{1, \ldots, N\}$ and for any $N \geq 2$. Hence in a game with *full sequential implementability*, the players' commitment levels are $\bar{\kappa} \triangleq 1/(N+1)$; this can be viewed as a *reference commitment*, against which one can compare the commitments achieved in non-trivial games. In a two-player game, for instance, the reference commitment is $\bar{\kappa} = 1/3$.

Lemma 5. *The commitment measure is invariant with respect to a positive-linear transformation of any agent's utility function.*

[10] If instead one considered $\max\{0, 1 - \alpha\varphi_i\}$ for any fixed $\alpha > 0$ as commitment measure, then the players' average commitment level would lie in the interval $[1 - (\alpha/N), 1]$ and would therefore tend to 1 as $N \to \infty$, completely eroding the informativeness of the commitment measure for large games.

Proof. For any player $i \in \mathcal{N}$, let $\hat{u}_i = \alpha_i u_i + \beta_i$, where $(\alpha_i, \beta_i) \in \mathbb{R}_{++} \times \mathbb{R}$. Then for any Nash equilibrium a^* of Γ the modified payoff ratios in Prop. 2 are the same as before:

$$\hat{q}_{ij}(a_i; a^*) = q_{ij}(a_i; a^*), \quad a_i \in \mathcal{A}_i, \quad i, j \in \mathcal{N},$$

so that Eqs. (10) and (11) together imply that the set $\mathcal{P}(a^*)$ remains unchanged. The claim now follows from Eqs. (13) and (14), concluding our proof. □

Invariance with respect to probabilistically equivalent representations of the players' payoffs is important because such rescaling does not affect any choice in the game. In particular, it does not affect the set of Nash equilibria and—as the preceding result shows—leaves the commitment measure unchanged as well.

Example 5. In the setting of the battle-of-the-sexes coordination game, let $a^* = (D, D)$ and $\hat{a}^* = (M, M)$ denote the two pure-strategy Nash equilibria. Using the results of Ex. 3 it is $\|\mathcal{P}(a^*)\| = \|\mathcal{P}(\hat{a}^*)\| = 1/4$. The expected flexibility is

$$\varphi_i(a^*) \triangleq \frac{\int_{\mathcal{P}(a^*)} p_i dp}{\|\mathcal{P}(a^*)\|},$$

so that

$$\varphi_1(a^*) = 4 \int_0^{1/2} (1 - 2p_1) p_1 dp_1 = \frac{1}{6}$$

for Ann, and

$$\varphi_2(a^*) = 2 \int_0^{1/2} (1 - 2p_1)^2 dp_1 = \frac{1}{3}$$

for Bert. Consequently, in the Nash equilibrium a^* favorable for Ann, by Eq. (14) her commitment measure is twice as large as Bert's (who merely achieves the reference commitment level of $\bar{\kappa} = 1/(N+1) = 1/3$ introduced in Ex. 4):

$$\kappa_1(a^*) = \frac{2}{3} > \frac{1}{3} = \kappa_2(a^*).$$

By symmetry, in the Nash equilibrium $\hat{a}^*$ favorable for Bert, one obtains $\kappa_i(\hat{a}^*) = \kappa_{-i}(a^*)$ for player $i \in \{1, 2\}$, so the expected commitment levels are reversed compared to a^*; see Fig. 3.

Example 6 (Entry Deterrence). Consider the classical two-player entry-deterrence game as, for example, in Rasmusen (2001, p. 94). The first player can either "enter" (E) or "not enter" ($\bar{E}$) the market. The second player, the incumbent monopolist, can either "accommodate" (A) or "fight" (F), e.g., by starting a price war. The extensive-form representation of the game and the players' payoffs are given in Fig. 4(a); the corresponding normal-form game Γ, depicted in Fig. 4(b), has two Nash equilibria: $a^* = (E, A)$ and $\hat{a}^* = (\bar{E}, F)$. Consider first $a^* = (E, A)$ and $p \in \Delta_2$. From the entrant's perspective this equilibrium can be sequentially implemented if and only if

$$U_1(\bar{E}, A; p) = 0 + 2p_1 + 0 \leq 2 = u_1(a^*) = U_1(a^*; p),$$

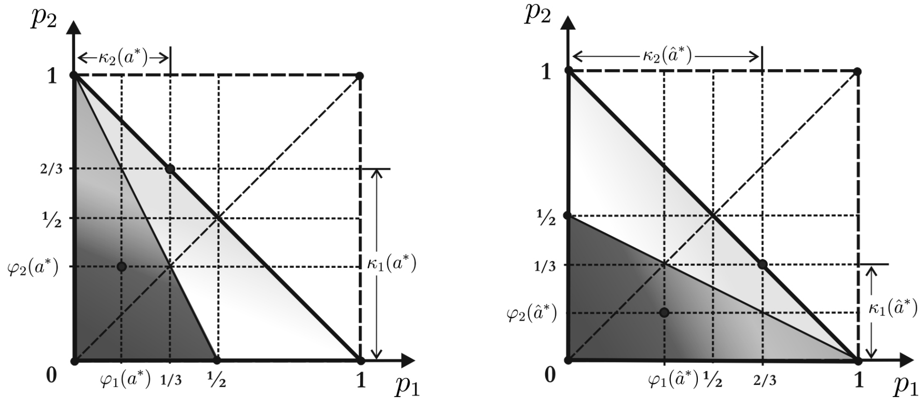

(a) (b)

Figure 3: Measures of commitment and flexibility in Ex. 5: (a) for $a^* = (D, D)$; (b) for $\hat{a}^* = (M, M)$.

so $\mathcal{P}_1(a^*) = \Delta_2$. On the other hand,

$$U_2(E, F; p) = -p_0 + p_2 + 2p_1 = -1 + 3p_1 + 2p_2 \leq 1 = u_2(a^*) = U_2(a^*; p),$$

whence $\mathcal{P}_2(a^*) = \{p \in \Delta_2 : p_2 \leq 1 - (3/2)p_1\} = \mathcal{P}(a^*)$, and

$$\|\mathcal{P}(a^*)\| = \int_0^1 \left(\int_0^{\max\{0, 1-(3/2)p_1\}} dp_2 \right) dp_1 = \int_0^{2/3} (1 - (3/2)p_1) dp_1 = \frac{1}{3}.$$

Determine now the expectations:

$$\int_{\mathcal{P}(a^*)} p_1 dp = \int_0^{2/3} (1 - (3/2)p_1) p_1 dp_1 = \frac{2}{27},$$

$$\int_{\mathcal{P}(a^*)} p_2 dp = \frac{1}{2} \int_0^{2/3} (1 - (3/2)p_1)^2 dp_1 = \frac{1}{9}.$$

The expected flexibilities are therefore $\varphi_1(a^*) = 2/9$ and $\varphi_2(a^*) = 1/3$, resulting in the commitment levels of $\kappa_1(a^*) = 5/9$ and $\kappa_2(a^*) = 1/3$. Hence, in the entry-accommodation Nash equilibrium, the entrant is 66.67% more committed than the incumbent.

Consider now the Nash equilibrium $\hat{a}^* = (\bar{E}, F)$, which contains a noncredible threat. For this, note that

$$U_2(A, \hat{a}_1^*; p) = 2p_0 + 2p_2 + p_1 = 2 - p_1 \leq 2 = u_2(\hat{a}^*) = U_2(\hat{a}^*; p),$$

for all $p \in \mathcal{P}_2(\hat{a}^*) = \Delta_2$. On the other hand,

$$U_1(E, \hat{a}_2^*) = -2p_0 + 0 + p_2 = -2 + 2p_1 + 3p_2 \leq 0 = u_1(\hat{a}^*) = U_1(\hat{a}^*; p),$$

if and only if $p \in \mathcal{P}_1(\hat{a}^*) = \{p : p_2 \leq (2/3)(1-p_1)\}$. Therefore, $\mathcal{P}(\hat{a}^*) = \mathcal{P}_1(\hat{a}^*) = \mathcal{P}_2(a^*) = \mathcal{P}(a^*)$. Hence, we obtain the expected flexibilities, $\varphi_1(a^*) = 1/3$ and $\varphi_2(a^*) = 2/9$, implying the commitments of $\kappa_1(a^*) = 1/3$ and $\kappa_2(a^*) = 5/9$. As a result, in the entry-deterrence Nash equilibrium $\hat{a}^*$ the situation is reversed, with the incumbent being 66.67% more committed than the entrant.

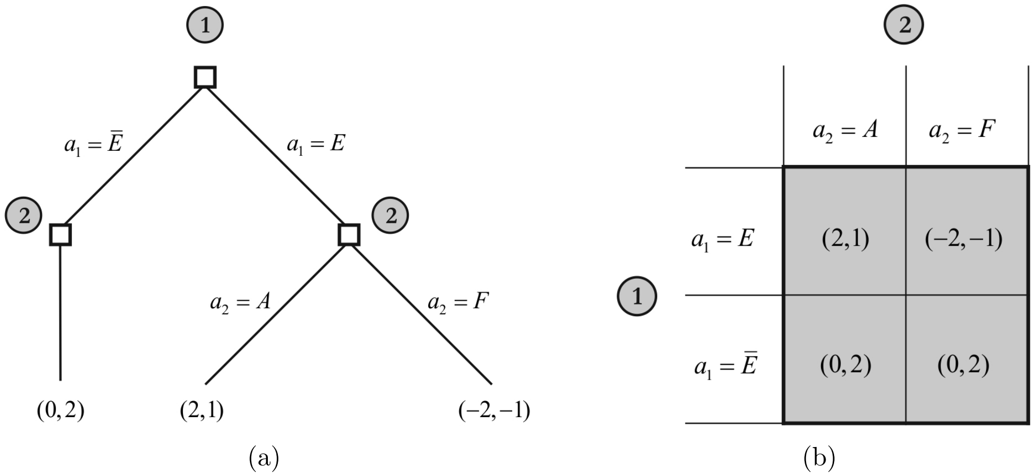

Figure 4: Extensive-form (a) and normal-form (b) representations of the dynamic entry-deterrence game in Ex. 6.

5 Conclusion

In the physical world, two events are considered “simultaneous” when they take place at exactly the same time in the same frame of reference. In a game, players’ actions are called “simultaneous moves” when, at the time one player chooses his action, he has not yet observed any of the other players’ actions. Simultaneity in a game is therefore usually viewed as an informational imperfection, rather than a condition brought about by the timing of actions. Yet, it is the timing of actions just as much as the resulting information structure that creates advantages through commitment (or lack thereof) when players move sequentially, each being able to observe earlier moves.

By viewing the Nash equilibria of a (finite) dynamic game as equilibria of the corresponding normal-form game Γ, we abstract from the particular structure and from any generically unique subgame-perfect equilibrium. Each Nash equilibrium a^* of Γ can be characterized by a set $\mathcal{P}(a^*)$ which contains vectors $p = (p_i)$ in a richer canonical extension $\hat{\Gamma}(p)$. The results in this paper embed the normal-form game Γ in the richer dynamic game $\hat{\Gamma}(p)$.[11] To determine player i’s commitment ability in a given Nash equilibrium a^* we first determine his expected flexibility $\varphi_i(a^*)$ as the expected value of p_i in a random instance $\hat{\Gamma}(p)$, conditional on the fact that a^* can be sequentially implemented (i.e., $p \in \mathcal{P}(a^*)$). That player’s commitment $\kappa_i(a^*) = \max\{0, 1 - N\varphi_i(a^*)\}$ is then the renormalized lack of flexibility, conditional on implementing the equilibrium in the canonical extension. Since κ_i can take on values between 0 and 1, and retains that range even on average across all players, for any size game, the measure retains informativeness. It is

[11] The somewhat unsurprising consequence of this embedding is that when the game $\hat{\Gamma}(p)$ is “sufficiently close” to Γ, that is, for $\|p\|$ small enough, any Nash equilibrium of Γ can be implemented sequentially in $\hat{\Gamma}(p)$. The implementation can be strong if all deviations from the best-responses in equilibrium lead to payoff differences that are uniformly bounded from below by some $\varepsilon > 0$. It is straightforward, and therefore asserted without proof, that the lack of “smooth” indifference around the equilibrium is not only sufficient but also necessary for a strong sequential implementation. The discreteness of choice around the equilibrium implies a certain robustness with respect to extensions of the game Γ, in the sense that the sets of possible equilibrium outcomes in Γ and $\hat{\Gamma}(p)$ coincide as long as $\|p\|$ is small enough.

invariant with respect to probabilistically equivalent representations of the players utility functions. In addition, the proposed measure of commitment is naturally adapted to the continuous nature of credibility (and commitment) in dynamic games, thus providing an indication about how small changes in the structure of the initial dynamic game can lead to a different subgame-perfect Nash equilibrium, without changing the underlying set of Nash equilibria.

Acknowledgements

The author is grateful for comments and suggestions by Bernhard von Stengel, several anonymous referees, as well as participants of the 2017 Conference on Game Theory and Management (GTM) at the Saint Petersburg State University and the 2017 European Meeting on Game Theory (SING13) at the Université Paris-Dauphine.

References

[1] Bennett, E., van Damme, E. (1990) "Demand Commitment Bargaining – The Case of Apex Games," CentER Discussion Paper No. 9062, Tilburg University, Tilburg, Netherlands.

[2] Berge, C. (1959) *Espaces Topologiques et Fonctions Multivoques.* Dunod, Paris. [English translation: *Topological Spaces*, Oliver and Boyd, Edinburgh, UK, 1963; reprinted: Dover Publications, Mineola, NY, 1997.]

[3] Breitmoser, Y. (2009) "Demand Commitments in Majority Bargaining or How Formateurs Get Their Way," *International Journal of Game Theory*, Vol. 38, No. 2, pp. 183–191.

[4] Caruana, G., Einav, L. (2008) "A Theory of Endogenous Commitment," *Review of Economic Studies*, Vol. 75, No. 1, pp. 99–116.

[5] Friedman, J.W. (1971) "A Non-Cooperative Equilibrium for Supergames," *Review of Economic Studies*, Vol. 38, No. 1, pp. 1–12.

[6] Fudenberg, D., Maskin, E. (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, Vol. 54, No. 3, pp. 533–554.

[7] Fudenberg, D., Tirole, J. (1991) *Game Theory*, MIT Press, Cambridge, MA.

[8] Gal-Or, E. (1985) "First Mover and Second Mover Advantages," *International Economic Review*, Vol. 26, No. 3, pp. 649–653.

[9] Henkel, J. (2002) "The 1.5th Mover Advantage," *RAND Journal of Economics*, Vol. 33, No. 1, pp. 156–170.

[10] Luce, R.D., Raiffa, H. (1957) *Games and Decisions*, Wiley, New York, NY.

[11] Mas-Colell, A., Whinston, M.D., Green, J.D. (1995) *Microeconomic Theory*, Oxford University Press, Oxford, UK.

[12] Montero, M., Vidal-Puga, J.J. (2007) "Demand Commitment in Legislative Bargaining," *American Political Science Review*, Vol. 101, No. 4, pp. 847–850.

[13] Nash, J.F. (1950) "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences*, Vol. 36, No. 1, pp. 48–49.

[14] Osborne, M.J., Rubinstein, A. (1984) *A Course in Game Theory*, MIT Press, Cambridge, MA.

[15] Selten, R. (1965) "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit," *Zeitschrift für die Gesamte Staatswissenschaft*, Vol. 121, pp. 301–324, 667–689.

[16] Selten, R. (1992) "A Demand Commitment Model of Coalition Bargaining," in: Selten, R. (Ed.) *Rational Interaction*, Springer, Berlin, Germany, pp. 245–282.

[17] von Stackelberg, H. (1934) *Marktform und Gleichgewicht*, Springer, Vienna, Austria.

[18] Weber, T.A. (2014) "A Continuum of Commitment," *Economics Letters*, Vol. 124, No. 1, pp. 67–73.

[19] Winter, E. (1994) "The Demand Commitment Bargaining and Snowballing Cooperation," *Economic Theory*, Vol. 4, No. 2, pp. 255–273.