Constrained Games: the impact of the attitude to adversary's constraints[◇]

Eitan Altman[*] and Eilon Solan[†]

***Abstract*—We consider properties of constrained games, where the strategy set available to a player depends on the choice of strategies made by other players. We show that the utilities of each player associated with that player's own performance and constraints are not sufficient to model a constrained game and to define equilibria; for the latter, one also needs to model how a player values the fact that other players meet their constraints. We study three different approaches to other players' constraints, and show that they exhibit completely different equilibrium behaviors. Further, we study a general class of stochastic games with partial information, and focus on the case where the players are indifferent to whether the constraints of other players hold.**

I. Introduction

Games with constraints have long been used for modeling and studying non-cooperative control in various areas [9], [11], [13]. Various models exist for constrained games; the simplest being one with orthogonal constraints, where the strategies of the players are restricted independently of each other [15]. A second model of interest is the model of Common Coupled Constraints (CCC) [15], [14], in which all players have a common convex non-orthogonal multi-strategy space. This model can be viewed as constraints that are *common to all users*. A unilateral deviation of a player from some feasible multi-strategy (one that satisfies the constraints) to another strategy that is feasible for that player, does not result, therefore, in the violation of constraints of other users. CCC have often been used in networking problems, where capacity constraints of links are naturally common. We study CCC in Section III.

In General Constrained Games (GCG) [8] the constraints are not necessarily common to all users. Therefore, if a single player deviates from a multi-strategy that is feasible for all players to another strategy that is feasible for the deviating player, the new multi-strategy need not be feasible for other users.

We argue that, in addition to the players' constraints, it is important to indicate the goal of each player with respect to the other players' constraints: does a player wish to prevent the constraints of another player to hold, or is the player indifferent to whether or not they hold. For example, when there are two players, and only one player, say, player 2 (P2), has constraints, the strategic behavior of P2 depends on the goals of player 1 (P1). If P1's primary goal is to prevent P2 to satisfy P2's constraints whenever possible, P2 must be very careful in choosing his strategy. If, on the other hand, P1's primary goal is to maximize his own payoffs, P2 has more strategies available. We show that qualitative aspects of the game differ when the players' attitudes towards the other players' constraints vary. We then provide, in Section V-B, a general equilibrium existence result for stochastic games in which the players are indifferent towards the other players' constraints.

II. The Model

Consider games with N players. The set of strategies available to each player i is $\mathcal{S}_i$ (which may be finite or infinite). Set $\mathcal{S} = \times_{i=1}^N \mathcal{S}_i$, and for every $x = (x_i)_{i=1}^N \in \mathcal{S}$ set $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_N)$, for every i. Given a multi-strategy of the other players, x_{-i}, the constraints allow player i to choose a strategy from the constrained strategy set $\mathcal{S}_i(x_{-i}) \subseteq \mathcal{S}_i$. A multi-strategy $x = (x_i)_{i=1}^N$ is *feasible* if $x_i \in \mathcal{S}_i(x_{-i})$ for every player i. The payoff[1] of player i is described by a function $u_i : \mathcal{S} \to \mathbb{R} \cup \{\pm\infty\}$, such that $u_i(x) \in \mathbb{R}$ for every feasible multi-strategy x. Thus, the payoff is defined even if the constraints are violated, but in this case the payoff may be infinite. These games are called General Constrained Games (GCG), since the constraints of one player may be different from those of other players. A feasible multi-strategy x is an *equilibrium* if $u_i(x) \geq u_i(x'_i, x_{-i})$, for every player i, and for every strategy $x'_i \in \mathcal{S}_i(x_{-i})$.

Example II.1. Power control in a cellular network
There are N mobile terminals, each of which has to send a transmission to a base station. Time is slotted, and at each time slot only one transmission can be successful. The strategic choice of each mobile i is the transmission power P_i that the mobile uses. The received power of mobile i is given by $h_i P_i$, where h_i is the channel gain, and it is assumed to be constant. At each time slot, the transmission of the mobile that maximizes the received power is successful. The goal of each mobile i is to minimize the transmission power P_i, subject to the constraint that its minimum expected throughput (which is the probability of a successful transmission) is at least some given bound. This is a GCG, since the success probability of a mobile depends on the actions of all other mobiles.

Dynamic version: Consider a dynamic situation, in which from time to time each mobile has to send transmissions to the base station, and the channel gains are not constant, but rather each one follows a Markov chain; that is, $(h_i(t))_i$ form N independent Markov chains, where $h_i(t)$ is the gain of mobile's i channel at time t. Assume also that there is a finite set of available power levels. Each mobile i knows at time t the number of transmissions waiting in the player's queue and the player's gain $h_i(t)$ at that time, but is unaware of the status of the other mobiles; all the player knows is the joint distribution of the number of transmissions they have to send and the joint distribution of their channel gains. The strategic choice of each mobile at every time slot t is its transmission power $P_i(t)$ at that time slot, and the received power from mobile i is the product $h_i(t)P_i(t)$. We may then consider the game where mobile i minimizes his own average power subject to his average expected throughput being at least some given bound. Alternative objectives will be discussed in Section V-B.

III. Games with common coupled constraints

The game has *Common Coupled Constraints* (CCC) if, for every multi-strategy $x = (x_k)_{k=1}^N$ and every pair of players i and j, $x_i \in \mathcal{S}_i(x_{-i}) \iff x_j \in \mathcal{S}_j(x_{-j})$. Thus, in games with CCC, for every multi-strategy, the constraints of one player are satisfied if and only if the constraints of all players are satisfied.

A *zero-sum* game is a two-player game in which $u_1 + u_2 = 0$. In this case we denote the payoff function of player 1 by U; that is, $u_1 = U$ and $u_2 = -U$. The *upper value* is $\inf_{y \in \mathcal{S}_2} \sup_{x \in \mathcal{S}_1(y)} U(x,y)$, and the *lower value* is $\sup_{x \in \mathcal{S}_1} \inf_{y \in \mathcal{S}_2(x)} U(x,y)$. The game has a *value* if the upper value and the lower value are the same.

[◇]This work was partially supported by the EuroNF Network of Excellence.
[*] INRIA B.P.93, 06902 Sophia-Antipolis Cedex, France.
[†] School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel.

[1]At present we do not require that the payoff functions or the functions $x_{-i} \mapsto \mathcal{S}_i(x_{-i})$ be measurable. When we state our results, we will indicate which conditions these functions should satisfy.

A feasible multi-strategy (x_*, y_*) is a *constrained saddle point* if

$$U(x_*, y_*) = \sup_{x \in \mathcal{S}_1(y_*)} U(x, y_*) = \inf_{y \in \mathcal{S}_2(x_*)} U(x_*, y).$$

We call $U(x_*, y_*)$ the *saddle point payoff*. In zero-sum games, the concepts of equilibrium and saddle point coincide. In unconstrained (finite) matrix games, as well as in (finite) matrix games with orthogonal constraints, a saddle point in mixed strategies always exists, and the saddle point payoff is the value of the game. Moreover, if (x_1, y_1) and (x_2, y_2) are two saddle points, then (x_1, y_2) and (x_2, y_1) are also saddle points. More generally, the following holds for zero-sum games (e.g., [7, p. 126]):

Lemma III.1. (*Minmax Theorem*)
Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be convex subsets of linear topological spaces, where $\mathcal{S}_2$ is compact. Consider a function $U : \mathcal{S}_1 \times \mathcal{S}_2 \to \mathbb{R}$ such that
– for each $x \in \mathcal{S}_1$, $y \to U(x, y)$ is convex and lower semi-continuous; and
– for each $y \in \mathcal{S}_2$, $x \to U(x, y)$ is concave.
Then there exists some $y_* \in \mathcal{S}_2$ such that

$$\inf_{y \in \mathcal{S}_2} \sup_{x \in \mathcal{S}_1} U(x, y) = \sup_{x \in \mathcal{S}_1} U(x, y_*) = \sup_{x \in \mathcal{S}_1} \inf_{y \in \mathcal{S}_2} U(x, y).$$

We conclude that under the conditions of the lemma, if $\mathcal{S}_1$ is compact as well then a saddle point exists. As we will see below, when the constraints are non-orthogonal the situation is completely different.

The following assumption will hold throughout the sequel.
Assumption A: For every $x \in \mathcal{S}_1$ and for every $y \in \mathcal{S}_2$ the sets $\mathcal{S}_1(y)$ and $\mathcal{S}_2(x)$ are nonempty.

In the context of games with CCC, this assumption is without loss of generality. Indeed, if, e.g., $\mathcal{S}_1(y) = \emptyset$, P2 cannot use the strategy y, since the constraints will be violated whatever P1 plays. Therefore, such strategies can be deleted from P2's strategy set.

Player II

		L	R
Pl. I	T	0 (2)	1 (0)
	B	−1 (0)	0 (0)

Fig. 1. Zero-sum game with coupled constraints

Example III.1. Matrix games.
In matrix games, the set of strategies of the players is the set of probability distributions over their respective (finite) sets of actions, the payoff function is multi-linear, and the constraints are multi-linear as well. Consider the constrained zero-sum matrix game that appears in Figure 1. The game is defined as follows:
(i) The strategies of P1 and P2 are the probability distributions written as row vectors: $x = (x(T), x(B))$ and $y = (y(T), y(B))$, respectively.
(ii) There are two matrices: U and D. The first corresponds to utilities and the second to constraints. The entries of U are the numbers given in the left-hand side of the corresponding boxes of the matrix in Figure 1. The entries of D correspond to the numbers appearing in parentheses in the right-hand side of each box in Figure 1.
(iii) P1 wishes to maximize the expected outcome $xUy^\intercal$, and P2 wishes to minimize it. ($y^\intercal$ is the transpose of y).
(iv) As in [15], the constraint is common to both players: $xDy^\intercal \le \rho$, where ρ is some constant, taken to be 0 in this example.

Consider the strategy $x_* = (1, 0)$ of P1 (choose T with probability 1). In order for the constraint to hold, P2 has to play $y_* = (0, 1)$ (choose R with probability 1), and the payoff is $x_* U y_*^\intercal = 1$. Since 1 is the maximal payoff we obtain

$$\max_x \min_{\{y : xDy^\intercal \le \rho\}} xUy^\intercal = U(T, R) = 1.$$

Next, assume that P2 chooses $y_* = (1, 0)$ (choose L with probability 1). To meet the constraint, P1 has to play $x_* = (0, 1)$ (choose B with probability 1), and the payoff is $x_* U y_*^\intercal = -1$. Since -1 is the minimal payoff we obtain

$$\min_y \max_{\{x : xDy^\intercal \le \rho\}} xUy^\intercal = U(B, L) = -1.$$

We conclude that the value does not exist. Moreover, we obtain the surprising unusual inequality

$$\max_x \min_{\{y : xDy^\intercal \le \rho\}} xUy^\intercal > \min_y \max_{\{x : xDy^\intercal \le \rho\}} xUy^\intercal. \qquad (1)$$

Observe that the two multi-strategies $x_* = (1, 0)$, $y_* = (0, 1)$ and $x_* = (0, 1)$, $y_* = (1, 0)$ are constrained saddle points, that yield different payoffs. In addition, $x_* = (1, 0)$, $y_* = (1, 0)$ and $x_* = (0, 1)$, $y_* = (0, 1)$ are not saddle points (the first one is not even feasible).

Example III.2. Networking games: parallel links. Consider a network with two parallel links connecting a source and a destination. Player i has a total demand ϕ^i, and has to decide how to split his demand between the two links. The strategy of player i is given by $x^i = (x_1^i, x_2^i)$, where x_ℓ^i is the amount of flow that player i sends over link ℓ. The capacity of link ℓ is C_ℓ units, and the cost per unit flow of link ℓ is $f_\ell(x_\ell)$, where $x_\ell = x_\ell^1 + x_\ell^2$ is the total flow on the link.

Consider an example with two players, $f_1 = 0$, $f_2(x_\ell) = x_\ell$, $C_1 = 5, C_2 = 10$, $\phi^1 = \phi^2 = 3$. Then the average cost for P1 is given by $J(\mathbf{x}) = x_2^1(x_2^1 + x_2^2)$. We assume that P1 wishes to minimize this cost and P2 wants to maximize it. P2 can be viewed as an intruder who wishes to degrade the performance of P1.

If P1 plays first, then P1 has a dominant strategy of shipping all his demands through link 1: $x_1^1 = 3, x_2^1 = 0$, and, independently of the strategy of P2, $J(\mathbf{x}) = 0$, so that

$$\min_{x^1} \max_{\{x^2 : (x^1, x^2) \text{ is feasible}\}} J(\mathbf{x}) = 0.$$

Similarly, if P2 plays first, P2 has the dominant strategy of sending all his flow through link 1. We then get

$$\max_{x^2} \min_{\{x^1 : (x^1, x^2) \text{ is feasible}\}} J(\mathbf{x}) = 1.$$

Thus, again we obtain the surprising inequality:

$$\min_{x^1} \max_{\{x^2 : (x^1, x^2) \text{ is feasible}\}} J(\mathbf{x}) < \max_{x^2} \min_{\{x^1 : (x^1, x^2) \text{ is feasible}\}} J(\mathbf{x}).$$

In this example there is a continuum of constrained saddle points: each multi-strategy x that satisfies $x_1^1 + x_1^2 = 3$ is a saddle point. Moreover, each saddle point yields a different payoff.

We now show that the surprising phenomenon that was exhibited in Examples III.1 and III.2 is common in games with common constraints.

We define the *unrestricted game* to be the one in which the constraints are relaxed in a way that the coupling between the constraints on each player is removed. The set of strategies of player 1 and 2 are then $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively.

Theorem III.1. *Consider a zero-sum game with CCC, and assume that in the unrestricted game the value exists (see Lemma III.1):*

$$\sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2} U(x,y) = \inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1} U(x,y). \tag{2}$$

Then the original constrained game satisfies

$$\sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y) \geq \inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1(y)} U(x,y). \tag{3}$$

Proof. Let v be the value of the unrestricted game:

$$v := \sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2} U(x,y) = \inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1} U(x,y).$$

For every $y \in \mathcal{S}_2$ one has $\mathcal{S}_1(y) \subseteq \mathcal{S}_1$, and therefore $\sup_{x\in\mathcal{S}_1(y)} U(x,y) \leq \sup_{x\in\mathcal{S}_1} U(x,y)$. Hence

$$\inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1(y)} U(x,y) \leq \inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1} U(x,y) = v.$$

By symmetry we obtain: $v \leq \sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y)$. Combining the two equations yields Eq. (3).

Remark III.1. **(i)** In zero-sum games with CCC, the maximization and minimization by players 1 and 2 respectively are restricted to feasible multi-strategies. The upper and lower values that appear in Eq. (3) are both taken over the feasible multi-strategies. This is in spite of the fact that in the left-hand side, the maximization of P1 is over all $\mathcal{S}_1$; P2 takes care that the constraints of P1 are satisfied. A symmetric argument holds for the right-hand side.
(ii) Eq. (3) holds also in constrained games without CCC. But it no longer has a useful interpretation, since (a) the lower value in the left-hand side is no longer restricted to multi-strategies that are feasible for P1, and (b) the upper value (in the right-hand side) is no longer restricted to multi-strategies that are feasible for P2.
(iii) The inequality $\inf_{y\in\mathcal{S}_2}\sup_{x\in\mathcal{S}_1} U(x,y) \geq \sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2} U(x,y)$, which represents the situation when no constraints are present, always holds. As Theorem III.1 states, in games with constraints in which the second mover must fulfill the constraints, the reverse inequality holds. This result is similar to the first mover advantage that is well recognized in economic theory. As we will see later, when the first mover has to play a strategy that ensures that the constraints are satisfied whatever the other player plays, this phenomenon no longer exists.

IV. Aggressive attitude to adversary's constraints

In this section we study zero-sum games with coupled constraints, in which each player's main goal is to prevent the other player from satisfying his constraints. We call this situation an "aggressive attitude to the adversary's constraints". (We assume however that whenever possible, a player will not violate his own constraints in order to prevent the constraints of the other player from being satisfied.)

The max-min value corresponds to the situation in which P1 moves first. Since the main goal of P2 is to prevent P1 from satisfying the constraints, P1 must choose a strategy that guarantees that the constraints are satisfied, whatever P2 plays. Let $\mathcal{G}_1 = \{x : x \in \mathcal{S}_1(y), \ \ \forall y \in \mathcal{S}_2\}$ be the set of those strategies of P1. Similarly, let $\mathcal{G}_2 = \{y : y \in \mathcal{S}_2(x), \ \ \forall x \in \mathcal{S}_1\}$ be the set of strategies of P2 that ensure that the constraints are satisfied, whatever P1 plays. The max-min value is given by $\sup_{x\in\mathcal{G}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y)$, while the min-max value is given by $\inf_{y\in\mathcal{G}_2}\sup_{x\in\mathcal{S}_1(y)} U(x,y)$.

We assume the following throughout the rest of the section.
Assumption B: $\mathcal{G}_1$ and $\mathcal{G}_2$ are nonempty sets.

As we will now show, in this setup the inequality in (2) in Theorem III.1 is reversed.

Theorem IV.1. *Consider a zero-sum game with GCG. Suppose that the value of the game in which each player i is restricted to strategies in $\mathcal{G}_i$ exists. When both players have an aggressive attitude to the adversary's constraints, then*

$$\sup_{x\in\mathcal{G}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y) \leq \inf_{y\in\mathcal{G}_2}\sup_{x\in\mathcal{S}_1(y)} U(x,y). \tag{4}$$

Remark IV.1. The reason that the inequality here is reversed is that a player P who plays first must be very cautious – that player must play a strategy in $\mathcal{G}_1$ or to the adversary's constraints, neither the value nor a saddle point may exist. In section V-B we show that if the players are indifferent to other players' constraints, an equilibrium always exists (and therefore in particular a saddle point in zero-sum games always exists).

Example IV.1. Consider the constrained game given in Figure 2. We

		Player II	
		L	R
Pl. I	T	1 (1)	0 (0)
	B	0 (0)	2 (1)

Fig. 2. Zero-sum game with constraints on player 2's strategies

		Player II	
		L	R
Pl. I	T	0 (1)	1 (0)
	B	2 (0)	0 (1)

Fig. 3. Zero-sum constrained game with two saddle points

assume that player 1 is not constrained, so that $\mathcal{G}_1 = \mathcal{S}_1$, whereas the constraints of player 2 are multi-linear: $\mathcal{S}_2(x) = \{y : xDy^\intercal \leq \rho\}$, where $\rho = 1/2$ is a constant.

We first show that the value need not exist. The set $\mathcal{G}_2$, which contains all the strategies of P2 that guarantee that P2 meets the constraints no matter what P1 does, is the singleton $\{(1/2, 1/2)\}$ (that is, the actions L and R have the same probability of being played). This implies that $\inf_{y\in\mathcal{G}_2}\sup_{x\in\mathcal{S}_1} U(x,y) = 1$ is obtained when P1 uses the action B with probability 1. On the other hand, $\sup_{x\in\mathcal{S}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y) = 3/4$. Indeed, P1 can get approximately $3/4$ by playing $(1/2+\varepsilon, 1/2-\varepsilon)$ where ε is a small positive number: the best response of P2 is $(1/2, 1/2)$, and the payoff is $3/4 - \varepsilon/2$. To see that P1 cannot guarantee more than $3/4$, observe that if P1 plays $(x(T), x(B))$, then, if $x(T) \leq 1/2$ the best response of P2 is L, which yields payoff $x(T) \leq 1/2$, whereas, if $x(T) \geq 1/2$, by playing $(1/2, 1/2)$ P2 ensures that the payoff is $1 - x(T)/2 \leq 3/4$.

We conclude that in this example,

$$\sup_{x\in\mathcal{G}_1}\inf_{y\in\mathcal{S}_2(x)} U(x,y) < \inf_{y\in\mathcal{G}_2}\sup_{x\in\mathcal{S}_1} U(x,y),$$

and the value does not exist.

We now argue that in this game a saddle point does not exist either. As mentioned above, $\mathcal{G}_2 = \{(1/2, 1/2)\}$, hence the only strategy of P2 that can be part of a saddle point is $y^* = (1/2, 1/2)$. However, P1's best reply to y^* is B, and P2's best reply to B is L and not y^*.

Though in general a saddle point need not exist, there are zero-sum games in which both players have an aggressive attitude to their adversary's constraints that do possess a saddle point, for example, a game in which the payoff function is constant. One may verify that, in Example III.1, if $\rho = 0$ then (T,R) is a saddle point.

Unlike games with CCC, when the players have an aggressive attitude towards the adversary's constraints, the payoff in all saddle points (if there are several saddle points) is the same.

Theorem IV.2. *Consider a zero-sum game with GCG, and assume that both players have aggressive attitude towards the other player's constraints. If (x_1, y_1) and (x_2, y_2) are two constrained saddle* *points, then $U(x_1, y_1) = U(x_2, y_2)$, and both (x_1, y_2) and (x_2, y_1) are saddle points.*

Proof: Since (x_1, y_1) and (x_2, y_2) are saddle points, one has $x_1, x_2 \in \mathcal{G}_1$ and $y_1, y_2 \in \mathcal{G}_2$. In particular, both (x_1, y_2) and (x_2, y_1) are feasible. Moreover,

$$U(x_1, y_1) = \inf_{y \in \mathcal{S}_2(x_1)} U(x_1, y), \quad U(x_2, y_2) = \sup_{x \in \mathcal{S}_1} U(x, y_2).$$

Since $y_2 \in \mathcal{G}_2 \subseteq \mathcal{S}_2(x_1)$,

$$U(x_1, y_1) = \inf_{y \in \mathcal{S}_2(x_1)} U(x_1, y) \leq U(x_1, y_2) \leq U(x_2, y_2).$$

By symmetry, we obtain $U(x_1, y_1) \geq U(x_2, y_2)$, so that $U(x_1, y_1) = U(x_1, y_2) = U(x_2, y_2)$. This implies

$$U(x_1, y_2) = U(x_1, y_1) = \inf_{y \in \mathcal{S}_2(x_1)} U(x_1, y)$$

$$\text{and} \quad U(x_1, y_2) = U(x_2, y_2) = \sup_{x \in \mathcal{S}_1(y_2)} U(x, y_2),$$

so that (x_1, y_2) is a saddle point. An analogous argument shows that (x_2, y_1) is a saddle point. ■

V. Indifference to opponents' constraints

We now assume that a player is not ready to suffer a loss in order to prevent the constraints of another player to hold. We call this behavior "indifference to opponents constraints". We first consider zero sum games, and provide a few illuminating examples. We then study non-zero sum stochastic games, and provide a general equilibrium existence result.

A. Zero-sum games

Example III.1 (continued) Consider first the matrix game in Figure 1, and interpret it as a game with a constraint only on P2, where P1 is indifferent to whether the constraints of P2 are satisfied or not. Since T is a dominant strategy for P1, the game has a saddle point (which is (T,R)).

Example IV.1 (continued) Consider the game in Figure 2. Suppose that P1 is not constrained: $\mathcal{S}_1(y) = \mathcal{S}_1$ for every y, and P2's constraints are defined with $\rho = 1/2$:

$$S_2(x) = \{y = (y(L), y(R)) \colon x(T)y(L) + x(B)y(R) \leq 1/2\}.$$

This gives

$$S_2(x) = \begin{cases} \{y = (y(L), y(R)) \colon y(L) \leq 1/2\} & \text{if } x(T) > 1/2, \\ \{y = (y(L), y(R)) \colon y(L) \geq 1/2\} & \text{if } x(T) < 1/2, \\ \{(1/2, 1/2)\} & \text{if } x(R) = 1/2. \end{cases}$$

Suppose that both players are indifferent to each other constraints. The unique equilibrium of the unconstrained game is $x(T) = 2/3$, $y(L) = 2/3$, which is not feasible. The best response for both players is given in Figure 4. The figure shows that a saddle point does not exist.

Example V.1. Consider the matrix game that appears in Figure 3, with $\rho = 0$. Both (T,R) and (B,L) are constrained saddle points, and the two saddle point payoffs differ. Moreover, (T,L) and (B,R) are not feasible, and in particular not constrained saddle points.

Matrix games are a special case of the general model we study in the next section, and therefore, as we show below, have a mixed equilibrium provided that the Slater condition holds.

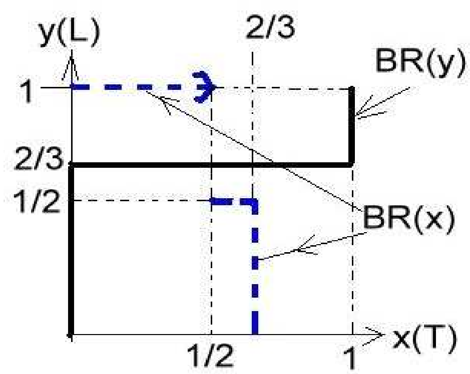

Fig. 4. The best responses in Example IV.1.

B. Stochastic non zero-sum constrained games

Below we provide an existence result for N-player stochastic games, which includes, as a special case, the static games discussed in previous sections (with the restriction that here we study games with finite action sets).

Related work. Zero-sum constrained stochastic games have been studied in the context where one player controls the transitions [5], [12], and in a setting where each player controls an independent Markov chain [10]. Non-zero-sum stochastic games with general constraints have been studied in [2], [3], [6]. Stochastic games with constraints on one side, where a player has an aggressive relation to his adversary's constraint, have been studied in [16].

Below, $\Delta(X)$ is the space of probability measures over the finite set X. Consider the following N-players game with a finite state space Σ. Each player i has:

- a finite set A_i of actions. Set $A := \prod_{i=1}^N A_i$;
- a finite set M_i of signals. Set $M := \prod_{i=1}^N M_i$;
- a stage payoff function $u_i : \Sigma \times A \times M \to \mathbb{R}$. Thus, the stage payoff depends on the current state, on the action profile, and on the signal profile;
- a discount factor $\lambda_i \in (0, 1)$;
- K constraint functions $d_i^k : \Sigma \times A \times M \to \mathbb{R}$, for $k = 1, ..., K$;
- for each constraint $k = 1, 2, ..., K$, a bound $\rho_i^k \in \mathbb{R}$ and a discount factor $\mu_i^k \in (0, 1)$.

For every stage $n \in \mathbb{N}$ there is a transition function $q_n : \Sigma \times A \times (M \times \Sigma \times A)^{n-1} \to \Delta(M \times \Sigma)$. Thus, given past play, nature chooses (i) the next state, and (ii) a signal profile, possibly in a correlated manner.

The game is played as follows. It starts at a given initial state $\zeta^1 \in \Sigma$. At each stage $n \geq 1$, the following happens: (i) Each player i chooses an action a_i^n. Set $a^n = (a_i^n)_{i=1}^N$. (ii) Nature chooses $(m^n, \zeta^{n+1}) \in \Sigma \times M$ according to $q(\bullet \mid \zeta^1, a^1, m^1, \zeta^2, a^2, m^2, \ldots, \zeta^n, a^n)$. (iii) Each player i is told of m_i^n, but is not told of the current state, the actions of the other players, or the payoffs. The signal, however, may contain some of this information.

The information available to player i at the beginning of stage n is the sequence of signals that the player has received so far: $m_i^1, ..., m_i^{n-1}$, and the sequence of actions that the player has chosen, $a_i^1, \ldots, a_1^{n-1}$. Hence, a behavior strategy of player i is a function:

$$s_i \;:\; \left(\bigcup_{n \in \mathbb{N}} (A_i \times M)^{n-1} \right) \to \Delta(A_i).$$

The discounted expected payoff of player i under a multi-strategy

$\mathbf{s} = (s_i)_{i=1}^N$ is[2]

$$\gamma_i(\mathbf{s}) := \mathbb{E}_{\zeta,\mathbf{s}}\left[\sum_{n\in\mathbb{N}}(1-\lambda_i)^{n-1}u_i(\zeta^n,a^n,m^n)\right]. \tag{5}$$

The discounted expected kth constraint of player i under the multi-strategy $\mathbf{s}$ is given by

$$\delta_i^k(\mathbf{s}) := \mathbb{E}_{\zeta,\mathbf{s}}\left[\sum_{n\in\mathbb{N}}(1-\mu_i^k)^{n-1}d_i^k(\zeta^n,a^n,m^n)\right]. \tag{6}$$

A strategy r_i for player i is *feasible* given a multi-strategy $\mathbf{s}_{-i}$ if $\delta_i^k(\mathbf{s}_{-i}, r_i) \leq \rho_i^k$. A multi-strategy $\mathbf{s} = (s_i)_i$ is feasible if s_i is feasible for player i given $\mathbf{s}_{-i}$, for every i. A feasible multi-strategy $\mathbf{s}$ is an *equilibrium* if $\gamma_i(\mathbf{s}) \geq \gamma_i(s_i', \mathbf{s}_{-i})$, for every player i and every strategy s_i' of player i which is feasible given $\mathbf{s}_{-i}$. Observe that a player's total payoff and total constraint are the discounted sum of stage payoffs and stage constraints. Such a case naturally occurs when the players have budget constraints, and the expenses of a player are the discounted sum of the player's stage expenses. Below we argue that our existence result holds in a model in which the total payoff and/or the total constraint are the (undiscounted) sum of stage payoffs or stage constraints, provided these summations are uniformly bounded.

We assume the following Slater condition holds.

Assumption C: For every multi-strategy $\mathbf{s}$, and every player i, there is a strategy r^i that strictly satisfies the constraints $d_i^k(r_i, \mathbf{s}_{-i}) < \rho_i^k, k = 1, \ldots, K$.

Theorem V.1. *Under Assumption C, in every discounted stochastic game with discounted constraints, an equilibrium exists.*[3]

It is worth noting that Example IV.1 that we have just studied is not a counter-example to Theorem V.1, since it does not satisfy Assumption C. Indeed, if $x(T) = x(B) = 1/2$, player 2 has no strategy y that satisfies that $xDy^\intercal < 1/2$.

Proof: A pure strategy of player i is a function

$$s_i : \left(\bigcup_{n\in\mathbb{N}}(A_i \times M)^{n-1}\right) \to A_i.$$

By Kuhn's theorem, every behavior strategy is equivalent to a mixture of pure strategies. Denote by Σ_i^P the space of pure strategies of player i, and by $\Sigma_i^B = \Delta(\Sigma_i^P)$ the space of his behavior strategies. Σ_i^P is a compact space in the product topology. Therefore Σ_i^B is compact in the weak topology, and it is convex.

Moreover, for every i and every multi-strategy $\mathbf{s}$, the functions $r_i \mapsto \gamma_i(\mathbf{s}_{-i}, r_i)$ and $r_i \mapsto \delta_i^k(\mathbf{s}_{-i}, r_i)$ are linear (and in particular continuous and quasi-concave).

For every player i and for every multi-strategy $\mathbf{s}$, define

$$F_i(\mathbf{s}) = \left\{r_i \in \Sigma_i^B : \delta_i^k(\mathbf{s}_{-i}, r_i) \leq \rho_i^k, \quad \forall k\right\}.$$

By Assumption C, this set is nonempty. Since $r_i \mapsto \delta_i^k(\mathbf{s}_{-i}, r_i)$ is linear, this set is convex and compact.

[2]Discounted games are an appropriate model when the interaction lasts many (possibly unknown number of) periods, and profits in the near future count more than profits in the far future. It is also appropriate when the players maximize the total sum of their stage payoffs (or minimize the total sum of their stage costs), and at every stage there is a fixed probability that the game terminates. We later argue that one can relax the termination condition.

[3]In our model each player observes a private signal. This implies that there is no common state variable, as the private history of the players differ. Moreover, without the knowledge of the strategies of the other players, a player cannot form a belief over the set of private histories that the other players observed. Therefore, previous work on constrained stochastic games [2], [6] cannot be applied here.

One can verify that the set-valued function $\mathbf{s} \mapsto F_i(\mathbf{s})$ is upper-semi-continuous. We prove that Assumption C ensures that it is lower-semi-continuous as well. Let $\mathbf{s}$ be arbitrary, and $r_i \in F_i(\mathbf{s})$. Let $(\mathbf{s}(l))_{l\in\mathbb{N}}$ be a sequence of strategies that converge to $\mathbf{s}$. We argue that there is a sequence $(r_i(l))_{l\in\mathbb{R}}$ that converges to r such that $r_i(l) \in \mathbf{s}(l)$ for each l. Indeed, by assumption C there is $\varepsilon > 0$, and a strategy $\widehat{s}_i$, such that $\delta_i^k(\mathbf{s}_{-i}, \widehat{s}_i) \leq \rho_i^k - \varepsilon$, for every k. Since δ_i^k is continuous,

$$\begin{aligned}\lim_{l\to\infty}\delta_i^k(\mathbf{s}_{-i}(l), r_i) &= \delta_i^k(\mathbf{s}_{-i}, r_i) \leq \rho_i^k,\\ \lim_{l\to\infty}\delta_i^k(\mathbf{s}_{-i}(l), \widehat{s}_i) &= \delta_i^k(\mathbf{s}_{-i}, \widehat{s}_i) \leq \rho_i^k - \varepsilon.\end{aligned}$$

Since δ_i^k is multi-linear, there is a sequence of numbers in the unit interval $(\alpha(l))_{l\in\mathbb{N}}$ that converges to 1 such that

$$\begin{aligned}&\delta_i^k(\mathbf{s}_{-i}(l), \alpha(l)r_i + (1-\alpha(l))\widehat{s}_i)\\ &= \alpha(l)\delta_i^k(\mathbf{s}_{-i}(l), r_i) + (1-\alpha(l))\delta_i^k(\mathbf{s}_{-i}(l), \widehat{s}_i) \leq \rho_i^k, \quad \forall k.\end{aligned}$$

The sequence of strategies $(s_i^*(l))_{l\in\mathbb{N}}$ that is defined by $s_i^*(l) := \alpha(l)r_i + (1-\alpha(l))\widehat{s}_i$ converges to r_i, and for each l one has $s_i^*(l) \in F_i(\mathbf{s}(l))$. Therefore F_i is lower-semi-continuous.

Define $G_i(\mathbf{s}) = \mathrm{argmax}_{r_i\in F_i(\mathbf{s})}\gamma_i(\mathbf{s}_{-i}, r_i)$. Then $G_i(\mathbf{s}) \subseteq \Sigma_i^B$ is convex and compact, and it has nonempty values. Since the set-valued function $\mathbf{s} \mapsto F_i(\mathbf{s})$ is both upper-semi-continuous and lower-semi-continuous, it follows that the set-valued function $\mathbf{s} \mapsto \times_{i=1}^N G_i(\mathbf{s})$ is upper-semi-continuous.

By Glicksberg's generalization of Kakutani's fixed point theorem, there is a fixed point $\mathbf{s}^* = (s_i^*)_{i=1}^N$: $s_i^* \in G_i(\mathbf{s}^*)$, for every i. The multi-strategy $\mathbf{s}_*$ is our desired equilibrium. ■

We notice that in Examples III.1 and III.2 all the saddle points are equilibria. If in Example III.1 we set $\rho = 1$, then there is a continuum of equilibria: each pair (x, y) such that $xDy^\intercal = 1$ is an equilibrium.

Total expected payoff. The proof of Theorem V.1 relies, in addition to Assumption C, on the following two properties of the model:

- The strategy spaces of the players are compact and convex sets in a metric space.
- The functions $r_i \mapsto \gamma_i(\mathbf{s}_{-i}, r_i)$, and $r_i \mapsto \delta_i^k(\mathbf{s}_{-i}, r_i)$ are continuous and quasi-concave.

A different model in which these properties are satisfied, and therefore the conclusion of Theorem V.1 holds, is the following. Suppose there is a "terminal state" $\zeta_0 \in \Sigma$, such that the game terminates once this state is reached. That is, the game stops at time τ, where $\tau := \min\{n \geq 1 : \zeta^n = \zeta_0\}$. Define the total expected payoff by

$$\widehat{\gamma}_i(\mathbf{s}) := \mathbb{E}_{\zeta,\mathbf{s}}\left[\sum_{n=1}^{\tau}u_i(\zeta^n,a^n,m^n)\right],$$

and the total expected k'th constraint by

$$\widehat{\delta}_i^k(\mathbf{s}) := \mathbb{E}_{\zeta,\mathbf{s}}\left[\sum_{n=1}^{\tau}d_i^k(\zeta^n,a^n,m^n)\right].$$

Since the summation in the definition of the functions γ_i and δ_i^k may be infinite, these functions may be undefined for some multi-strategies. A simple sufficient conditions that ensures that these functions are well defined is that $\mathbb{E}_{\zeta,\mathbf{s}}[\tau] < \infty$ for every pure stationary multi-strategy in the problem with full information [12]. For other sufficient conditions see [1].

The functions $\widehat{\gamma}_i$ and $\widehat{\delta}_i^k$ are multi-linear, and in particular continuous and quasi-concave. It follows that in a stochastic game in which the total reward of each player i is either γ_i or $\widehat{\gamma}_i$, and his total k'th constraint is either δ_i^k or $\widehat{\delta}_i^k$, an equilibrium exists (provided that $\widehat{\gamma}_i(\mathbf{s})$ and $\widehat{\delta}_i(\mathbf{s})$ are well defined for every i and every $\mathbf{s}$).

As an application example, consider the dynamic setting in Example II.1. Assume that the gains $h_i(t)$ are unknown to the mobiles. Periodically, the base station broadcasts a pilot signal with a known power; the power of the pilot received by mobile i is a private signal that allows the mobile to estimate the current value of $h_i(t)$. Assume also that each mobile has a battery with a finite amount of energy. After each transmission, the remaining energy depletes by the amount of energy used for transmission. The objective of a mobile is to maximize the total expected lifetime of its battery, subject to a constraint on the success probability at each slot (which should be larger than some constant). This problem can be modeled within the framework of this section.

VI. Conclusions

We have considered games with various types of constraints, and observed phenomena that are new with respect to unconstrained games. We have shown that for each constraint on a given player i, one has to define how the other players value the violation of the constraint of player i. In zero-sum games, depending on the attitude to other's constraints, the value need not exist, and surprisingly the max-min can be larger than the min-max. Other variations in which different players have different attitudes to other players' constraints can also be studied. We finally studied non-zero-sum stochastic games with constraints, and identified conditions for the existence of an equilibrium.

References

[1] E. Altman, *Constrained Markov Decision Processes*, Chapman and Hall/CRC, 1999.

[2] E. Altman and A. Shwartz, "Constrained Markov games: Nash equilibria", *Annals of Dynamic Games*, vol. 5, Birkhäuser, V. Gaitsgory, J. Filar and K. Mizukami, editors, 213-221, 2000.

[3] E. Altman, K. Avrachenkov, N. Bonneau, M. Debbah, R. El-Azouzi and D. Sadoc Menasche, "Constrained Cost-Coupled Stochastic Games with Independent State Processes" *Operations Research Letters*, Vol 36 pp. 160-164, 2008.

[4] E. Altman and Y. Hayel, "A Stochastic Evolutionary Game Approach to Energy Management in a Distributed Aloha Network", IEEE INFOCOM, April 2008.

[5] E. Altman, S. Sarkar and E. Solan, "Constrained Single-Controller Stochastic Games", SMCtools, Nantes, France, Oct. 2007.

[6] J. Alvarez-Mena and O. Hernandez-Lerma, "Existence of nash equilibria for constrained stochastic games", *Mathematical Methods in Operations Research*, 63:261–285, 2006.

[7] J. P. Aubin, *Optima and Equilibria, An Introduction to Nonlinear Analysis*, Springer-Verlag, 1993.

[8] G. Debreu, "A social equilibrium existence theorem", *Proceedings of the National Academy of Sciences of the U.S.A*, 38:886-893, 1952.

[9] R. El-Azouzi and E. Altman, "Constrained traffic equilibrium in routing", *IEEE Transaction on Automatic Control*, 48:1656-1660, 2003.

[10] E. Gómez-Ramírez, K. Najim and A.S. Poznyak, "Saddle-point calculation for constrained finite Markov chains". *Journal of Economic Dynamics and Control*, 27:1833-1853, 2003.

[11] M. T. Hsiao and A. A. Lazar. "Optimal decentralized flow control of Markovian queueing networks with multiple controllers", *Performance Evaluation*, 13:181-204, 1991.

[12] L. C. M. Kallenberg (1983), *Linear Programming and Finite Markovian Control Problems*, Mathematical Centre Tracts 148, Amsterdam.

[13] T. Larsson and M. Patriksson. "Side constrained traffic equilibrium models – traffic management through link tolls". In P. Marcotte and S. Nguyen, editors, *Equilibrium and Advanced Transportation Modelling*, pages 125–151. Kluwer Academic Publishers, 1998.

[14] L. Pavel, "An extension of duality to a game-theoretic framework", *Automatica*, Volume 43 , Issue 2, pp. 226-237, 2007.

[15] J. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games", *Econometrica*, 33:520-534, 1965.

[16] N. Shimkin, "Stochastic games with average cost constraints," Annals of the International Society of Dynamic Games, Vol. 1, Eds. T. Basar and A. Haurie, Birkhauser, 1994.