

Nash Equilibria And Partition Functions Of Games With Many Dependent Players

Elisabeth Kraus, TQA Research Group, University of Munich.

Simon Lentner¹, Algebra and Number Theory, University of Hamburg.

Abstract

We discuss and solve a model for a game with many players, where a subset of truly deciding players is embedded into a hierarchy of dependent agents.

These interdependencies modify the game matrix and the Nash equilibria for the deciding players. In a concrete example, we recognize the partition function of the Ising model and for high dependency we observe a phase transition to a new Nash equilibrium, which is the Pareto-efficient outcome.

An example we have in mind is the game theory for major shareholders in a stock market, where intermediate companies decide according to a majority vote of their owners and compete for the final profit. In our model, these interdependency eventually forces cooperation.

Contents

1	Introduction	2
2	Definition Of The Hierarchical Game	3
2.1	Conditional Influences	6
2.2	Payoff Mechanisms	7
3	Hierarchical Games On Trees Are Ising Models	8
4	Solution Of The Hierarchical Game	9
4.1	Steps For The Solution	9
4.2	Example To Step 1: Easy Hierarchical Graph	10
4.3	Example To Step 2-4: Prisoner's Dilemma	11
5	Open Questions	13

¹Corresponding author: simon.lentner@uni-hamburg.de

Acknowledgements

Both authors would like to thank Prof. Martin Schottenloher and the research group TQA at the LMU Munich for hospitality and support. The second author is supported by the DFG RTG 1670, University of Hamburg.

1 Introduction

Roughly 20 years ago, an exciting new set of methods has been introduced into game theory: An underlying game G , such as the Minority Game, played by a large ensemble of players, can be analyzed and often solved using methods from statistical physics. For comprehensive overviews on the development of this subject, see [CMZ05] or [Coolen05]. The analysis exhibits critical points where phase transitions appear in the thermodynamic limit of many players and it provides models for the emergence of mutual cooperation. The evolution of a large set of agents with prescribed strategies or learning mechanisms have been studied as dynamical systems e.g. in [CZ97], [Coolen05] or [NM92].

In this article, we want to focus on the influence of a large ensemble of agents on the game theory of a given game G , including game matrix and Nash equilibria. In the model we study, the game G is embedded into a hierarchy of automata/agents, that pass decisions by majority votes. The top of the hierarchy remains a set of new active/deciding players, which now play a new transformed game Γ . We then wish to understand how the game theory of Γ compared to G changes, depending on the given intermediate hierarchy. This has been solved by the first author as part of her diploma thesis [Kraus11].

More specifically, in section 2 we suppose that we are given a weighted directed graph \mathcal{H} and a game $G = \langle L, S, u \rangle$ played by a subset of the vertices \mathcal{H}_V of \mathcal{H} called *executive players* $L \subset \mathcal{H}_V$. We define a transformed game $\Gamma = \langle \Lambda, \Sigma, \nu \rangle$ with new players $\Lambda \subset \mathcal{H}_V$. The graph hereby is imagined as a hierarchy of agents with executive players L at the bottom of the hierarchy, while new *deciding players* Λ at the top of the hierarchy successively control the behaviour of the agents according to their influence. Conversely, the *payoff of G* for all executive players L is finally collected by the deciding players according to the natural bargaining process in this situation (Shapely value). Other payoff mechanisms are possible and got discussed in [Kraus11] as well.

The main example for this model we have in mind is the stock market, where the deciding players Λ send instructions through a graph that represents the structure of the mutual ownerships of the companies. Finally, some executive companies L play a executive game G in “reality”, such as prisoner’s dilemma or minority game. So we ask, how the stock market and mutual ownerships of the intermediate

agents/companies have altered the game G to Γ .

To solve Γ thermodynamically, we need more than a partition function summing over all possible strategies as done e.g. in [Coolen05]. Rather, the conditional probabilities of one agent's decision influencing another one have to be calculated. They correspond physically to k -point correlation functions (see definition 8) and especially for $\Lambda = \emptyset$ (no deciding players) the overall expression reduces to the partition function.


In section 3, we recognize that for treelike hierarchy graphs our k -point correlation functions coincide with a generalized Ising model on the graph \mathcal{H} . This enables us in principle to write down the game matrix, Nash equilibria and phase transitions of the game Γ for a given game G whenever the Ising model for \mathcal{H} is accessible. Of particular interest to us is the case where \mathcal{H} is a random graph, which has been solved in [DGH10].

In section 4 we demonstrate the approach and methods developed in this article. We solve and thoroughly analyse an example of an executive prisoners dilemma G being transformed to a hierarchical game Γ with again two deciding players, but with a certain hierarchy of agents between deciders and executive G -players.

Especially, we can establish a phase transition in the game Γ , if the branching factor of the tree is sufficiently high (otherwise we get only a tipping point) as the mutual influence approaches a critical threshold. The phase transition in Γ separates a phase with the (defecting) Nash equilibrium in G from a phase corresponding to the (cooperating) Pareto-efficient outcome. Roughly spoken, if the mutual dependency in decision making gets high, egoistic strategies become unstable and mutual cooperation emerges.

2 Definition Of The Hierarchical Game

In the following we suppose to be given a weighted directed graph \mathcal{H} , whose vertices contain among others executive players L playing a game G . We suppose G to have only two moves. The graph should be imagined as a hierarchy with executive players L at the bottom of the hierarchy. The key notion of this article is then a transformed game $\Gamma = \langle \Lambda, \Sigma, \nu \rangle$ with new deciding players Λ at the top of the hierarchy, who successively control the behaviour of the agents and collect the G -payoff according to their influence.



Definition 1 (Executive Game G). From now on, let $G = \langle L, \{S_i\}_{i \in L}, \{u_i\}_{i \in L} \rangle$ be a game in normal form with players $L = \{1 \dots n\}$ and each player $i \in L$ having two strategies $S_i = S = \{\pm 1\}$. The overall strategy set is hence $S^L = \prod_{i \in L} S_i = \{\pm 1\}^n$ and we denote the payoff for each player $i \in L$ by $u_i : S^L \rightarrow \mathbb{R}$.

We denote by A^B the set of all maps between A and B and by $\mathbb{R}A$ the vector space spanned by the set A .

Definition 2 (Hierarchy Graph \mathcal{H}). Let $\mathcal{H} = (\mathcal{H}_V, \mathcal{H}_E, \{f_{vw}\}_{vw \in \mathcal{H}_E})$ be a connected, directed, weighted graph with vertex set \mathcal{H}_V and directed edges $vw \in \mathcal{H}_E$ with positive weights $f_{vw} > 0$ for $v, w \in \mathcal{H}_V$.

We denote the direct predecessors and successors of vertices $w, v \in \mathcal{H}_V$ by

$$\text{pre}(w) = \{v \in \mathcal{H}_V \mid vw \in \mathcal{H}_E\} \quad \text{suc}(v) = \{w \in \mathcal{H}_V \mid vw \in \mathcal{H}_E\}.$$

We further denote by $\mathcal{H}_0 \subset \mathcal{H}_V$ all vertices without predecessors and without loss of generality we assume the predecessor weights to be normed:

$$\forall w \in \mathcal{H}_V \setminus \mathcal{H}_0 : \sum_{v \in \text{pre}(w)} f_{vw} = 1$$

Definition 3. (Hierarchical Game $\Gamma = \mathcal{H}G$) Suppose a fixed game $G = \langle L, S, u \rangle$ and a fixed hierarchy graph \mathcal{H} with $L \subset \mathcal{H}_V$. The transformed hierarchical game $\Gamma = \mathcal{H}G := \langle \Lambda, \Sigma, \nu \rangle$ consists of

- A set of deciding players $\Lambda := \mathcal{H}_0 = \{\lambda_1 \dots \lambda_m\}$.
- A strategy set $\Sigma_\lambda = \Sigma = S^L$ for each deciding player $\lambda \in \Lambda$. Such a strategy formulates a G -strategy-command to each executive player $i \in L$. The overall strategy set is hence Σ^Λ .
- A payoff function $\nu_\lambda : \Sigma^\Lambda \rightarrow \mathbb{R}$ for each deciding player $\lambda \in \Lambda$ given by

$$\nu_\lambda = \left(\sum_{i \in L} \phi_\lambda^{(i)} \cdot u_i \right) \circ \pi \circ P_{L|\Lambda}. \quad (1)$$

The function $\pi : \mathbb{R}\Sigma^L = \mathbb{R}(S^L)^L \rightarrow \mathbb{R}S^L$ is given by restricting a set of G -strategy-commands for each executive player $(\sigma_i^{(j)})_{i,j \in L} \in \Sigma^L$ to the strategies chosen for the respective player $(\sigma_i^{(i)})_{i \in L} \in S^L$.

The functions $P_{L|\Lambda} : \mathbb{R}\Sigma^\Lambda \rightarrow \mathbb{R}\Sigma^L$ and $\phi_\lambda : \Lambda \rightarrow \mathbb{R}^L$ depending on the hierarchy \mathcal{H} will be defined in what follows:

- $P_{B|A} : \mathbb{R}\Sigma^A \rightarrow \mathbb{R}\Sigma^B$ for subsets $A, B \subset \Sigma$ denotes the conditional influence of players A on players B and should be read as a $(|A| + |B|)$ -point-function. A deciding player $\lambda \in \Lambda$ has been defined to have a strategy

$$\sigma_\lambda = (\sigma_\lambda^{(i)})_{i \in L} \in S^L =: \Sigma$$

formulating the aim to have each executive player i using strategy $\sigma_\lambda^{(i)}$. These G -strategy-commands $\sigma = (\sigma_\lambda)_{\lambda \in \Lambda} \in \Sigma$ of all deciding players λ compete along the hierarchy graph and determine an overall outcome probability distribution $P_{L|\Lambda}(\sigma) \in \mathbb{R}\Sigma^L$ as described in the next section.

- $\phi_\lambda : \Lambda \rightarrow \mathbb{R}^L$ describes, how much of the payoff earned by each of the executive player $i \in L$ can be finally collected by a deciding player λ . The condition $\sum_{\lambda \in \Lambda} \phi_\lambda^{(i)} = 1$ is needed. In [Kraus11] we have discussed different payoff collection mechanisms, but in the following we will restrict ourselves to the natural result of a bargaining process between the deciding players Λ determined by the the Shapely value ([OR94]). This particular choice has moreover the nice property to only depend on the conditional influences $P_{L|\Lambda}$.

Remark 4. Stock Market

An easy application of this model is a stock market. The game $\Gamma = \langle \Lambda, \Sigma, \nu \rangle$ is played by deciding players Λ (e.g. major stockholders). The graph represents mutual ownerships of companies that pass the instructions of the deciding players via (2) to the executive players: These equations represent a voting in each node, that weights the possessions of the direct predecessors (respective direct owners) together with a small percentage D , the free float of randomly voting minor stockholders. The executive players L play the game $G = \langle L, S, u \rangle$ according to the instructions they get - they act as agents and aren't players in a game theoretical sense. The payoff which the executive players get is returned to the deciding players weighted by the Shapely value: The more influence deciding player λ has on executive player i , the more is λ getting of i 's payoff u_i .

So we ask how the stock market and mutual ownerships of the intermediate agents/companies have altered the game G to Γ . Roughly we find that if the mutual dependency in decision making gets high, egoistic strategies become unstable and mutual cooperation emerges.

2.1 Conditional Influences

We yet have to explain the function $P_{L|\Lambda} : \mathbb{R}\Sigma^\Lambda \rightarrow \mathbb{R}\Sigma^L$. First consider a neighbourhood graph $\mathcal{H}_{p,P}$ consisting of a point p with predecessors P . Suppose a *yes-no*-decision process, where $\sigma_v^{(p)} \in \{\pm 1\}$ represents commands of each $v \in P$ to p . The process shall be a vote in p , where every predecessor v has votes according to the weight f_{vp} and a percentage of $D \in]0, 1[$ votes randomly $\sim \mathcal{N}(0, \sigma_{\mathcal{N}}^2)$.

Lemma 5 (Single Vote). *For the neighbourhood graph $\mathcal{H}_{p,P}$ the probability for a result $+1$ in the point p under some given condition $(\sigma_v^{(p)})_{v \in P}$ is*

$$P_{p|P}^{single} \left(\sigma_v^{(p)} = +1 \mid (\sigma_v^{(p)})_{v \in P} \right) = 1 - P_{\mathcal{N}}(0, \sigma_{\mathcal{N}}^2)(-C) \approx \frac{1}{2} - \frac{1}{2} \tanh(-aC)$$

at which $C = \frac{1-D}{D} \sum_{v \in P} f_{vp} \sigma_v^{(p)}$ and $a = \sqrt{\frac{2}{\pi \sigma_{\mathcal{N}}^2}}$. (2)

Proof. Denote by $X^{(p)} \sim \mathcal{N}(0, \sigma_{\mathcal{N}}^2)$ the Gaussian random variable of the random voters.

$$\begin{aligned} P_{p|P}^{single} \left(\sigma_v^{(p)} = +1 \mid (\sigma_v^{(p)})_{v \in P} \right) &= P_{p|P}^{single} \left(DX^{(p)} + (1-D) \sum_{v \in P} f_{vp} \sigma_v^{(p)} \geq 0 \right) \\ &= P_{p|P}^{single} \left(X^{(p)} \geq - \underbrace{\frac{1-D}{D} \sum_{v \in P} f_{vp} \sigma_v^{(p)}}_C \right) \\ &= 1 - P_{\mathcal{N}}(0, \sigma_{\mathcal{N}}^2)(-C) \\ &\approx \frac{1}{2} - \frac{1}{2} \tanh(-aC) \text{ with } a = \sqrt{\frac{2}{\pi \sigma_{\mathcal{N}}^2}} \end{aligned}$$

For approximation we use the similarity of the normal distribution and the tangens hyperbolicus (see [Kraus11]). □

Provided that \mathcal{H} does not contain directed cycles, the entire voting process goes on iteratively and we obtain straight-forward by induction for any $A \subset \mathcal{H}_0, B \subset \mathcal{H}$:

$$P_{B|A}(\sigma) = \sum_{\tau \in \{\pm 1\}^{\mathcal{H}_V}, \tau|_A = \sigma} \tau|_B \cdot \prod_{p \in \mathcal{H}_V \setminus \mathcal{H}_0} P_{p|\text{pre}(p)}^{single} \left(\tau|_p \mid \tau|_{\text{pre}(p)} \right)$$

If \mathcal{H} does contain directed cycles then there is no terminating voting process. We nevertheless propose in complete analogy to statistical mechanics to assign in such a situation the conditional probabilities, which clearly reduce to the previous expression when no directed cycle is present:

$$P_{B|A}(\sigma) = \frac{1}{Z_{B|A}} \sum_{\tau \in \{\pm 1\}^{\mathcal{H}_V}, \tau|_A = \sigma} \tau|_B \cdot \prod_{p \in \mathcal{H}_V \setminus \mathcal{H}_0} P_{p|\text{pre}(p)}^{single} \left(\tau|_p \mid \tau|_{\text{pre}(p)} \right)$$

with the following now nontrivial normalization constant called partition function

$$Z_{B|A}(\sigma) := \sum_{\tau \in \{\pm 1\}^{\mathcal{H}_V}, \tau|_A = \sigma} \prod_{p \in \mathcal{H}_V \setminus \mathcal{H}_0} P_{p|\text{pre}(p)}^{\text{single}} \left(\tau|_p \mid \tau|_{\text{pre}(p)} \right).$$

This expression can be justified by a random experiment as follows: Let the probability space be $\Omega := \{\pm 1\}^{\mathcal{H}_V}$ with product measure

$$P(\tau) := \prod_{p \in \mathcal{H}_V \setminus \mathcal{H}_0} P_{p|\text{pre}(p)}^{\text{single}} \left(\tau|_p \mid \tau|_{\text{pre}(p)} \right).$$

Take as events $\Omega_A(\sigma) \subset \Omega$ to be the event that holds $\tau|_A = \sigma$ and analogously for $\Omega_B(\sigma')$. Then the conditional probability for $\Omega_B(\sigma')$ under the condition $\Omega_A(\sigma)$ is defined as

$$P_{B|A} = \frac{P(\Omega_B(\sigma') \cap \Omega_A(\sigma))}{P(\Omega_A(\sigma))}.$$

Plugging in the product measure $P(\tau)$ and taking a formal linear combination over the outcome $\sigma' = \tau|_B$ yields the formula above.

2.2 Payoff Mechanisms

Once the conditional probabilities $P_{L|\Lambda}(\sigma) \in \mathbb{R}\Sigma^L$ for given strategies σ_λ of each deciding player $\lambda \in \Lambda$ have been evaluated, this determines the behaviour of the executive players to

$$\tau = (\pi \circ P_{L|\Lambda})(\sigma) \in \mathbb{R}S^L.$$

This strategy produces in the game G a payoff $u_i(\tau)$ for each executive player $i \in L$. So how is this payoff collected finally by the executive players in Λ ? Denote by $\phi_\lambda^{(i)}$ the percentage of payoff of executive player i that is collected by deciding player λ , the *payoff collecting mechanism*. Then the overall payoff function of the game Γ is

$$\begin{aligned} \nu_\lambda : \Sigma^\Lambda &\longrightarrow \mathbb{R} \\ \nu_\lambda(\sigma) &= \left(\sum_{i \in L} \phi_\lambda^{(i)} \cdot u_i \right) \circ \pi \circ P_{L|\Lambda}(\sigma). \end{aligned} \quad (3)$$

Example 6 (Payoff by shares). *The most intuitive payoff collecting mechanism for treelike hierarchy graphs is payoff proportional to the amount of shares of the executive player i indirectly held by deciding player λ :*

$$\phi_\lambda^{(i)} = \sum_{\text{paths from } \lambda \text{ to } i} \prod_{vw \in \text{path}} f_{vw}$$

This fulfills $\sum_{\lambda \in \Lambda} \phi_\lambda^{(i)} = 1 \forall i$ (see [Kraus11]). An example for this payoff collecting mechanism is the paying of dividends proportional to the amount of stocks held by an owner.

However, in the following we will restrict ourselves to the natural result of a bargai-

ning process between the deciding players Λ according to their influence - the payoff is hence determined by the Shapely value [OR94] of the obvious coalition function:

$$z_i : 2^{|\Lambda|} \longrightarrow \mathbb{R}$$

$$K \longmapsto \frac{P(s_i = 1 | \sigma_\lambda^{(i)} = 1 \ \forall \lambda \in K, \sigma_\lambda^{(i)} = -1 \ \forall \lambda \notin K) - P(s_i = 1 | r_\lambda^{(i)} = -1 \ \forall \lambda)}{2P(s_i = 1 | r_\lambda^{(i)} = 1 \ \forall \lambda) - 1}$$

This particular choice for a payoff collection mechanism has moreover the nice property to only depend on the conditional influences $P_{L|\Lambda}$.

Lemma 7. *For the payoff function by Shapely value we get*

$$\phi_\lambda^{(i)} = \sum_{K \subseteq \Lambda, \lambda \in K} \frac{(|K| - 1)! (|\Lambda| - |K|)!}{|\Lambda|} (z_i(K) - z_i(K \setminus \{\lambda\})). \quad (4)$$

Proof. The necessary scaling condition $\sum_{\lambda \in \Lambda} \phi_\lambda^{(i)} = 1$ is fulfilled (see [OR94]). \square

3 Hierarchical Games On Trees Are Ising Models

In the following section we prove that the conditional influence can be calculated by using an isomorphic Ising model. For the common definition of the Ising model, see for example [Nolting07]. Some first analogies are obvious:

Ising model	hierarchical game
particles in a graph \mathcal{H}	same graph \mathcal{H}
spin of particle v : $\sigma_v \in \{\pm 1\}$	strategy of $v \in \mathcal{H}_V$: $\sigma_v = (\sigma_v^{(1)} \dots \sigma_v^{(n)}) \in \{\pm 1\}^n$
interaction J_{vw}	(with D modified) weights $f_{vw} \cdot \frac{1-D}{D}$
external magnetic field (is set to 0)	some systematic bias (not treated)

In addition to that, k -point-functions are needed for conditional probabilities.

Definition 8. *Let A, B be disjoint subsets of the nodes in the Ising model. Let the nodes beyond A , named $N(A)$, be the nodes of $\mathcal{H}_V \setminus (A \cup B)$, that fulfill the following condition: Every path to any $v' \in B$ hits at least one $v \in A$. With nodes beyond B (named $N(B)$) defined analogously, the nodes between A and B are $N(A, B) := \mathcal{H}_V \setminus (N(A) \cup N(B))$.*

With $k = |A| + |B|$ the k -point-function is

$$\{\pm 1\}^{|B|} \times \{\pm 1\}^{|A|} \longrightarrow \mathbb{R}$$

$$(\sigma', \sigma) \longmapsto \langle \sigma' \mid N(A, B) \mid \sigma \rangle$$

$$= \sum_{\sigma', \sigma \text{ fixed}} \exp(-\beta H_{N(A, B)}) \quad (5)$$

where $H_{N(A, B)}$ is the Hamiltonian function of the restricted graph just containing the nodes $N(A, B)$ and $\beta = \frac{1}{k_B T}$ is the inverse temperature in the Ising model.

Lemma 9. Let $\sigma_v \in \{\pm 1\}$ be the spin of particle v , $\sigma \in \{\pm 1\}^{|A|}$ the spins of particles in A and let the external magnetic field be 0. Then the conditional probability for σ_v given σ is

$$P(\sigma_v = 1 \mid \sigma) = \frac{\langle \sigma_v = 1 \mid N_{A, \{v\}} \mid \sigma \rangle}{\sum_{\sigma_v} \langle \sigma_v \mid N_{A, \{v\}} \mid \sigma \rangle} \quad (6)$$

Proof. See [Kraus11]. □

Remark 10. The Ising model without external magnetic field and with constant interaction J is exactly solvable in one dimension, see [Nolting07]. In this case, the conditional probability for two particles v, v' with spins σ, σ' and distance a is

$$P(\sigma' \mid \sigma) = \frac{\cosh^a(\beta J) \pm \sinh^a(\beta J)}{2 \cosh^a(\beta J)} \quad (7)$$

where the case "+" occurs when $\sigma = \sigma'$ and "-" if $\sigma \neq \sigma'$. This remark will be needed for the example in subsection 4.2.

Theorem 11. Every hierarchical game on a graph (as defined in section 2) that fulfills the condition, that $\forall i \in L$ the restricted graph of the nodes $N_{\Lambda, i}$ is a tree, is isomorphic to an Ising model such that a process in the game that leads from fixed $\sigma_\lambda^{(i)} \forall \lambda \in \Lambda$ to the strategy $\sigma_i^{(i)} = s_i$ of one $i \in L$ is equivalent to a process in an Ising model on the same graph with interactions $J_{vw} = f_{vw} \frac{1-D}{D}$, inverse temperature $\beta = \sqrt{\frac{2}{\pi \sigma_\lambda^2}}$ and no external magnetic field.

Isomorphic hereby means that the conditional influence and the k -point-functions coincide.

Proof. For details, see section 4 in [Kraus11].

As the graph had to be specified to a tree-like graph, the process fixates the strategies step by step. Therefore is it enough to look at the local fixation of a strategy in one node and to compare (2) and (6). Doing this gives the interactions and the inverse temperature as mentioned above. □

Remark 12. The restriction to tree-like graphs seems harsh on the first sight. However even a one-dimensional model shows interesting behaviour (see subsection 4.2) and also the calculation of Ising models on random graphs as done in [DGH10] typically requires the graph to be at least locally treelike.

4 Solution Of The Hierarchical Game

4.1 Steps For The Solution

Let $\Gamma = \langle \Lambda, \Sigma, \nu \rangle$ be a hierarchical game as defined in section 2 with $\Lambda = \{\lambda_1, \lambda_2\}$ and $L = \{1, 2\}$. (The restriction to two players allows the use of payoff matrices.)

1. Determining the conditional influence

With two players in each Λ and L , the conditional influence depends on only

four variables obtained from the Ising model and equation (6).

$$\begin{aligned}
\text{conditional influence on 1 : } x &= P_{L|\Lambda}(s_1 = 1 \mid \sigma_{\lambda_1}^{(1)} = -1, \sigma_{\lambda_2}^{(1)} = 1) \\
y &= P_{L|\Lambda}(s_1 = 1 \mid \sigma_{\lambda_1}^{(1)} = 1, \sigma_{\lambda_2}^{(1)} = 1) \\
\text{conditional influence on 2 : } \bar{x} &= P_{L|\Lambda}(s_2 = 1 \mid \sigma_{\lambda_1}^{(2)} = 1, \sigma_{\lambda_2}^{(2)} = -1) \\
\bar{y} &= P_{L|\Lambda}(s_2 = 1 \mid \sigma_{\lambda_1}^{(2)} = 1, \sigma_{\lambda_2}^{(2)} = 1)
\end{aligned}$$

2. Building up the pre-payoff matrix

The pre-payoff matrix assigns to every combination of strategies $\sigma = (\sigma_{\lambda_1}, \sigma_{\lambda_2})$ in Σ^Λ the expected payoff u_i to executive player $i \in L$ in game G :

$$\begin{aligned}
\mathbb{E}[(u_1, u_2) \mid \sigma] &= (u_1, u_2) \circ \pi \circ P_{L|\Lambda}(\sigma) \\
&= \sum_{s \in S} P_{L|\Lambda}(s \mid \sigma_{\lambda_1}, \sigma_{\lambda_2}) \cdot u(s) \\
&= \sum_{s_1} \sum_{s_2} P_{L|\Lambda}(s_1 \mid \sigma_{\lambda_1}^{(1)}, \sigma_{\lambda_2}^{(1)}) \cdot P_{L|\Lambda}(s_2 \mid \sigma_{\lambda_1}^{(2)}, \sigma_{\lambda_2}^{(2)}) \cdot u(s_1, s_2)
\end{aligned}$$

3. Building up the payoff matrix

To get the payoff matrix ν_λ for Γ as defined in equation (3) multiply every item in the pre-payoff matrix with

$$\begin{pmatrix} \phi_{\lambda_1}^{(1)} & \phi_{\lambda_2}^{(1)} \\ \phi_{\lambda_1}^{(2)} & \phi_{\lambda_2}^{(2)} \end{pmatrix} = \begin{pmatrix} \frac{y-x}{2y-1} & \frac{\bar{x}+\bar{y}-1}{2\bar{y}-1} \\ \frac{x+\bar{y}-1}{2y-1} & \frac{\bar{y}-x}{2\bar{y}-1} \end{pmatrix}.$$

4. Usual methods

Now that there is a payoff matrix for the game, the usual game theoretical methods can be applied to find Nash equilibria and phase transitions.

4.2 Example To Step 1: Easy Hierarchical Graph




Figure 1: One-dimensional hierarchical game

In the one-dimensional hierarchical game, the deciding and executive players are connected by chains of a, b, c or d edges. Therefore the weights are 1, except down at the executive players (weights $\frac{1}{2}$). Let D be $\frac{1}{2}$. The conditional influence can now be calculated with equation (7) which leads to

$$x = P_{L|\Lambda}(s_1 = 1 \mid \sigma_{\lambda_1}^{(1)} = -1, \sigma_{\lambda_2}^{(1)} = 1) = \frac{\mu(-, a) \cdot \mu(+, c)}{\mu(-, a) \cdot \mu(+, c) + \mu(+, a) \cdot \mu(-, c)}$$

$$y = P_{L|\Lambda}(s_1 = 1 \mid \sigma_{\lambda_1}^{(1)} = 1, \sigma_{\lambda_2}^{(1)} = 1) = \frac{\mu(+, a) \cdot \mu(+, c)}{\mu(+, a) \cdot \mu(+, c) + \mu(-, a) \cdot \mu(-, c)}$$

where $\mu(\pm, k) = \cosh^{k-1}(\beta) \cosh(\frac{\beta}{2}) \pm \sinh^{k-1}(\beta) \sinh(\frac{\beta}{2})$.

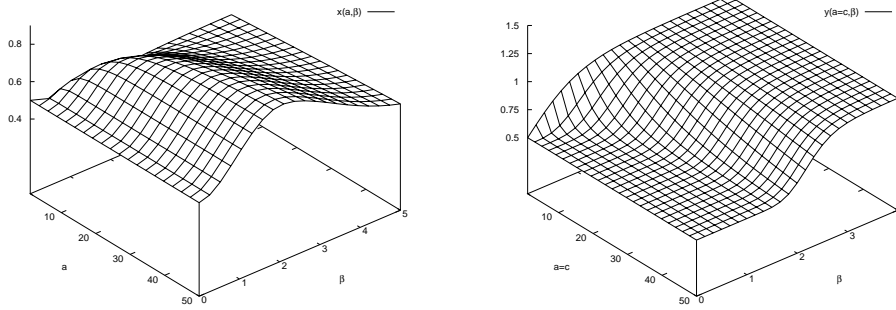


Figure 2: Conditional influence x with fixed c (left) and conditional influence y with $a = c$ (right)

As x is a measure for player 1 obeying rather λ_2 than λ_1 if their instructions differ, x is growing if a increases and the influence of λ_1 therefore decreases as it can be seen on the left. However, for small β (which means high temperature T) x is close to $\frac{1}{2}$ no matter how far the deciding players are from each other. As y shows how much player 1 is likely to obey λ_1 and λ_2 if they agree, the right graph shows how y is close to 1 if both deciding players are near to 1 at a low temperature. If the distance and the temperature increase, 1 tends to choose its strategy randomly with probability $\frac{1}{2}$.

4.3 Example To Step 2-4: Prisoner's Dilemma

Let the game G be the well-known prisoner's dilemma with payoff matrix

	$C_{(\text{operation})}$	$D_{(\text{efection})}$
C	(1, 1)	(-3, 3)
D	(3, -3)	(-1, -1)

Let the conditional influence be symmetric, so it goes down to just two variables x and y :

$$x = P_{L|\Lambda}(s_1 = C \mid \sigma_{\lambda_1}^{(1)} = D, \sigma_{\lambda_2}^{(1)} = C) = \bar{x}$$

$$y = P_{L|\Lambda}(s_1 = C \mid \sigma_{\lambda_1}^{(1)} = C, \sigma_{\lambda_2}^{(1)} = C) = \bar{y}$$

The complete payoff matrix has been calculated in [Kraus11].

Depending on x and y the hierarchical game Γ is isomorphic to one of the following games with unique Nash equilibrium $\hat{\sigma}$ and the tipping points for these three states are $x = \frac{2-y}{3}$ and $x = \frac{y+1}{3}$:

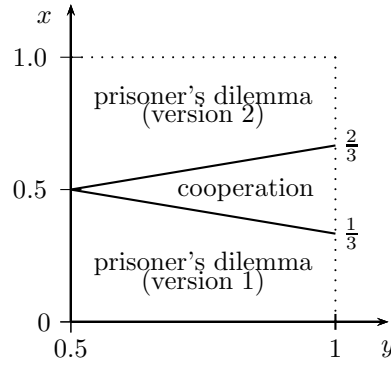


Figure 3: Illustration of the situation of Γ depending on x and y

1. Γ is a **prisoner's dilemma** where

for λ_1 : $(C, C) \simeq$ Cooperation, $(D, C) \simeq$ Defection

for λ_2 : $(C, C) \simeq$ Cooperation, $(C, D) \simeq$ Defection

That means, λ_1 identifies with 1 and λ_2 identifies with 2.

Hence the unique Nash equilibrium is $\hat{\sigma} = (\hat{\sigma}_{\lambda_1}, \hat{\sigma}_{\lambda_2}) = ((D, C), (C, D))$.

2. Γ is a **prisoner's dilemma** where

for λ_1 : $(C, C) \simeq$ Cooperation, $(C, D) \simeq$ Defection

for λ_2 : $(C, C) \simeq$ Cooperation, $(D, C) \simeq$ Defection

That means, λ_1 identifies with 2 and λ_2 identifies with 1.

Hence the unique Nash equilibrium is $\hat{\sigma} = ((C, D), (D, C))$.

3. **Cooperation:**

For λ_1 and λ_2 , the strategy (C, C) dominates every other strategy and therefore $\hat{\sigma} = ((C, C), (C, C))$.

Because of the symmetry the payoffs of λ_1 and λ_2 coincide: $\nu_{\lambda_1}^{x,y}(\hat{\sigma}) = \nu_{\lambda_2}^{x,y}(\hat{\sigma})$. Hence the value of the game, i.e. the payoff in the Nash equilibrium $\hat{\sigma}$, is as follows:

$$\nu_{\lambda_1}^{x,y}(\hat{\sigma}) = \begin{cases} -1 + 2x & x < \frac{2-y}{3} \\ -1 + 2y & \frac{2-y}{3} < x < \frac{y+1}{3} \\ 1 - 2x & x > \frac{y+1}{3} \end{cases}$$

Remark 13. *In the 1-dimensional hierarchy considered in section 4.2 we had*

$$x = \frac{\mu(-, a) \cdot \mu(+, c)}{\mu(-, a) \cdot \mu(+, c) + \mu(+, a) \cdot \mu(-, c)}$$

$$y = \frac{\mu(+, a) \cdot \mu(+, c)}{\mu(+, a) \cdot \mu(+, c) + \mu(-, a) \cdot \mu(-, c)}$$

where $\mu(\pm, k) = \cosh^{k-1}(\beta) \cosh(\frac{\beta}{2}) \pm \sinh^{k-1}(\beta) \sinh(\frac{\beta}{2})$ and the inverse temperature $\beta = \sqrt{\frac{2}{\pi \sigma_{\mathcal{N}}^2}}$ depended on the random minority voters.

Hence in the one-dimensional hierarchy we get the tipping points above, but x and y are still smooth functions in β . On the contrary, for a two-dimensional hierarchy, x and y would exhibit proper non-analytical phase transitions in the thermodynamic limit, turning the tipping points into proper phase transitions.

5 Open Questions

Question 14. *If we choose the Shapely value as payoff mechanism as above, the overall transformed game Γ depends only on the game G and the correlators. What can be said in general about the game theory of Γ compared to G without explicit knowledge of the correlators (under some reasonable, general assumptions)?*

Question 15. *It would be interesting to derive closed expressions for the correlators of an Ising model on a locally treelike random graph, similarly to the partition functions obtained in this case in [DGH10]; it is to be expected that e.g. the 2-point correlator depends only on the distance. This would yield a very nice explicitly solvable model with phase transition for games on randomly dependent agents.*

Question 16. *Our model does not necessarily require the graph to be a directed tree, see end of section 2.1. In fact, mutual dependencies might be more realistic. Then the following issues arise:*

- *Even in the easiest case, the partition sum does (to our surprise) not coincide with the partition sum of the Ising model. Rather, there are corrections for every directed loop. It would be nice to explain this behaviour and/or derive expressions for the partition sum, phase transition etc. in this modified versions using the same techniques from statistical physics as for the Ising model (transfer matrix for small dimension, mean field method for large dimension resp. branching number).*

- Alternatively, one might introduce a relaxation time, so the model gains a time dependence. This could be interesting to study non-stationary behaviour.

Question 17. *Can there be obtained statistical real-world evidence (and quantified), that the existence of inter-dependency on the path between deciding players and actual decision (as modelled in this article) increases cooperation?*

References

- [Kraus11] E. Kraus: *Anwendung des Ising-Modells auf den Aktienmarkt*. Diplomarbeit (2011)
- [DGH10] S. Dommers, C. Giardinà, R. van der Hofstad: *Ising models on power-law random graphs*. *Journal of Statistical Physics*, 141(4):638-660 (2010)
- [CZ97] D. Challet, Y.-C. Zhang: *Emergence of cooperation and organization in an evolutionary game*. *Physica A: Statistical Mechanics and its Applications*. Volume 246, 3-4, p. 407-418 (1997)
- [CMZ05] D. Challet, M. Marsili, Y.-C. Zhang: *Minority Games*. University Oxford Press (2005)
- [Coolen05] ACC Coolen: *The Mathematical theory of minority games - statistical mechanics of interacting agents*. Oxford University Press (2005)
- [OR94] M. Osborne, A. Rubinstein: *A Course in Game Theory*. MfT (1994). Chapter 14.4.
- [Nolting07] W. Nolting: *Grundkurs Theoretische Physik 6: Statistische Physik*. Springer (2007)
- [NM92] M. A. Nowak, R. M. May: *Evolutionary games and spatial chaos*. *Nature* (1992)