



Incentive Alignment of Business Processes

Tobias Heindel^(✉)  and Ingo Weber 

Chair of Software and Business Engineering, Technische Universitaet Berlin, Berlin,
Germany

{heindel,ingo.weber}@tu-berlin.de

Abstract. Many definitions of business processes refer to business goals, value creation, profits, etc. Nevertheless, the focus of formal methods research on business processes lies on the correctness of the execution semantics of models w.r.t. properties like deadlock freedom, liveness, or completion guarantees. However, the question of whether participants are *interested* in working towards completion – or in participating in the process at all – has not been addressed as of yet.

In this work, we investigate whether inter-organizational business processes give participants incentives for achieving the business goals: in short, whether *incentives are aligned* within the process. In particular, fair behavior should pay off and efficient completion of tasks should be rewarded. We propose a game-theoretic approach that relies on algorithms for solving stochastic games from the machine learning community. We describe a method for checking incentive alignment of process models with utility annotations for tasks, which can be used for *a priori* analysis of inter-organizational business processes. Last but not least, we show that the soundness property is a special case of incentive alignment.

Keywords: Inter-organizational business processes · Incentive alignment · Collaboration · Choreography

1 Introduction

Many definitions of what a business process is refer to business goals [29] or value creation [7], but whether process participants are actually incentivized to contribute to a process has not been addressed as yet. For *intra*-organizational processes, this question is less relevant; motivation to contribute is often based on loyalty, bonuses if the organization performs well, or simply that tasks in a process are part of one's job. Instead, economic modeling of intra-organizational processes often focuses on cost, e.g. in activity-based costing [12], which can be assessed using model checking tools [9] or simulation [5].

For *inter*-organizational business processes, such indirect motivation cannot be assumed. A prime example of misaligned incentives was the \$2.5B write-off in Cisco's supply chain in April 2001 [20]: success of the overall supply chain was grossly misaligned with the incentives of individual participants. (This happened despite the availability of several game theoretic approaches for analyzing

incentive structures for the case of supply chains [4].) Furthermore, modeling incentives accurately is actually possible in cross-organizational processes, e.g., based on contracts and agreed-upon prices. With the advent of *blockchain* technology [30], it is possible to execute cross-organizational business processes or choreographies as smart contracts [18,28]. The blockchain serves as a neutral, participant-independent computational infrastructure, and as such enables collaboration across organizations even in situations characterized by a *lack of trust* between participants [28]. However, as there is no central role for oversight, it is important that incentives are properly designed in such situations, e.g., to avoid unintended –possibly devastating– results, like those encountered by Cisco. In fact, a main goal of the Ethereum blockchain is, according to its founder Vitalik Buterin, to create “a better world by aligning incentives”¹.

In this paper, we present a framework for incentive alignment of inter-organizational business processes based on game theory. We consider BPMN models with suitable annotation concerning the utility² of activities, very much in the spirit of activity-based costing (ABC) [12, Chapter 5]. In short, fair behavior should pay off and participants should be rewarded for efficient completion of process instances. In more detail, we shall consider BPMN models as stochastic games [24] and formalize incentive alignment as “good” equilibria of the resulting game. Which equilibria are the desirable ones depends on the business goals w.r.t. which we want align incentives. In the present paper, we focus on *proper completion* and *liveness* of activities. Interestingly, the soundness property [2] will be rediscovered as the special case of incentive alignment within a single organization that rewards completion of every activity.

The overall contribution of the paper is a framework for incentive alignment of business process models, particularly in inter-organizational settings. Our approach is based on game theory and inspired by advances on the solution of stochastic games from the machine learning community, which has developed algorithms for the practical computation of Nash [22] and correlated equilibria [16,17]. The framework focuses on checking incentive alignment as an *a priori* analysis of business processes specified as BPMN models with activity-based utility annotations. Specifically, we:

1. describe a principled method for translating BPMN-models with activity-based costs to stochastic games [24]
2. propose a notion of incentive alignment that we prove to be a conservative extension of Van der Aalst’s soundness property [2],
3. illustrate the approach with a simplified order-to-cash (O2C) process.

We pick up the idea of incentive alignment for supply chains [4] and set out to apply it in the realm of inter-organizational business processes. From a technical point of view, we are interested in extending the model checking tools for cost analysis [9] for BPMN process models to proper collaborations, which we model as stochastic games [24]. This is analogous to how the model checker

¹ <https://www.ikiguide.com/ethereum/>, accessed 8-3-2020.

² We shall use utility functions in the sense of von Neumann and Morgenstern [19].

PRISM has been extended from Markov decision processes to games [14]. We keep the connection with established concepts from the business process management community by showing that incentive alignment is a conservative extension of the soundness property (see Theorem 1). Our approach hinges on algorithms [16, 22] for solving the underlying stochastic games of BPMN process models, which are sufficient for checking incentive alignment.

The remainder of the paper is structured as follows. We introduce concepts and notations in Sect. 2. On this basis, we formulate two versions of incentive alignment in Sect. 3. Finally, we draw conclusions in Sect. 4. The proof of the main theorem can be found in the extended version [8].

2 Game Theoretic Concepts and the Petri Net Tool Chest

We now introduce the prerequisite concepts for stochastic games [24] and elementary net systems [23]. The main benefit of using a game theoretic approach is a short list of candidate definitions of equilibrium, which make precise the idea of a “good strategy” for rational actors that compete as players of a game. We shall require the following two properties of an equilibrium: (1) no player can benefit from unilateral deviation from the “agreed” strategy and (2) players have the possibility to base their moves on information from a single (trusted) mediator. The specific instance that we shall use are *correlated equilibria* [3, 10] as studied by Solan and Vieille [25].³ We take ample space to review the latter two concepts, followed by a short summary of the background on Petri nets.

We use the following basic concepts and notation. The cardinality and the powerset of a set M are denoted by $|M|$ and $\wp M$, respectively. The set of real numbers is denoted by \mathbb{R} and $[0, 1] \subseteq \mathbb{R}$ is the unit interval. A probability distribution over a finite or countably infinite set M is a function $p: M \rightarrow [0, 1]$ whose values are non-negative and sum up to 1, in symbols $\sum_{m \in M} p(m) = 1$. The set of all probability distributions over a set M is denoted by $\Delta(M)$.

2.1 Stochastic Games, Strategies, Equilibria

We proceed by reviewing core concepts and central results for stochastic games [24], introducing notation alongside; we shall use examples to illustrate the most important concepts. The presentation is intended to be self-contained such that no additional references should be necessary. However, the interested reader might want to consult standard references or additional material, e.g., textbooks [15, 21], handbook articles [11], and surveys [26]. We start with the central notion.

Definition 1 (Stochastic game). *A stochastic game G is a quintuple $G = \langle N, S, A, q, u \rangle$ that consists of*

³ Nash equilibria are a special case, which however have drawbacks that motivate Aumann’s work on the more general correlated equilibria [3, 10].

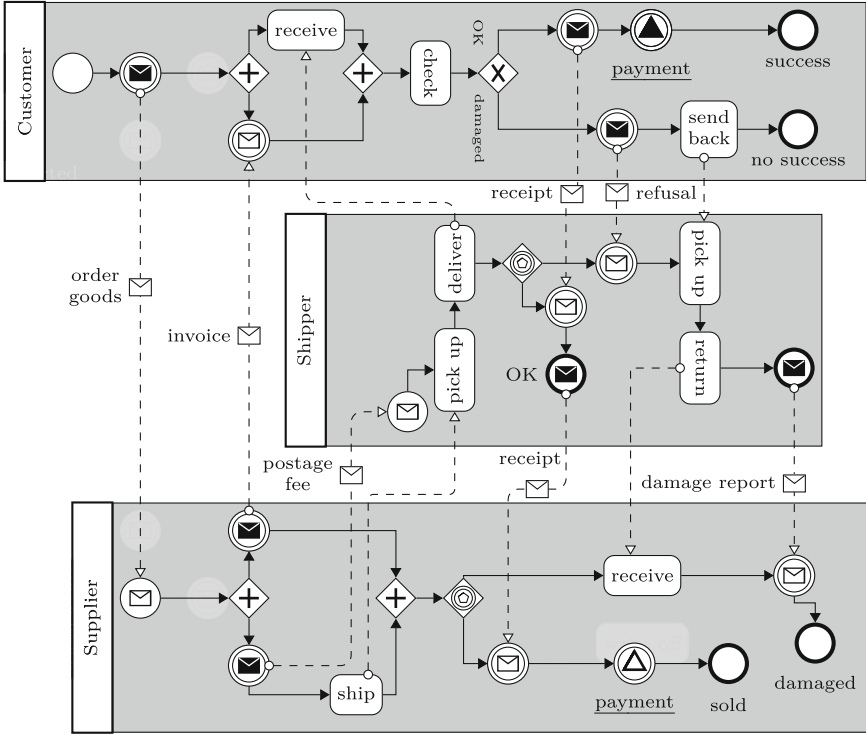


Fig. 1. A simplified order-to-cash process

- a finite set of players $N = \{1, \dots, |N|\}$ (ranged over by i, j, i_n , etc.);
- a finite set of states S (ranged over by s, s', s_n , etc.);
- a finite, non-empty set of action profiles $A = \prod_{i=1}^{|N|} A^i$ (ranged over by a, a_n , etc.), which is the Cartesian product of a player-indexed family $\{A^i\}_{i \in N}$ of sets A^i , each of which contains the actions of the respective player (ranged over by a^i, a_n^i , etc.);
- a non-empty set of available actions $A^i(s) \subseteq A^i$, for each state $s \in S$ and player i ;
- probability distributions $q(\cdot \mid s, a) \in \Delta(S)$, for each state $s \in S$ and every action profile $a \in A$, which map each state $s' \in S$ to $q(s' \mid s, a)$, the transition probability from state s to state s' under the action profile a ; and
- the payoff vectors $u(s, a) = \langle u^1(s, a), \dots, u^{|N|}(s, a) \rangle$, for each state $s \in S$ and every action profile $a = \langle a^1, \dots, a^{|N|} \rangle \in A$.

Note that players always have some action(s) available, possibly just a dedicated idle action, see e.g. [13].

The BPMN model of Fig. 1 can be understood as a stochastic game played by a shipper, a customer, and a supplier. Abstracting from data, precise timings, and similar semantic aspects, a state of the game is a state of an instance of

the process, which is represented as a token marking of the BPMN model. The actions of each player are the activities and events in the respective pool, e.g., the *ship* task, which *Supplier* performs after receiving an order from the *Customer* and payment of the postage fee to *Shipper*. Action profiles are combinations of actions that can (or must) be executed concurrently. For example, sending the order and receiving the order after the start of the collaboration may be performed synchronously (e.g., via telephone). The available actions of a player in a given state are the tasks or events in the respective pool that can be executed or happen next – plus the idle action. The transition probabilities for available actions in this BPMN process are all 1, such that if players choose to execute certain tasks next, they will be able to do so if the chosen activities are actually available actions. As a consequence, all other transition probabilities are 0.

One important piece of information that we have to add to a BPMN model via annotations is the utility of tasks and events. In analogy to the ABC method, which attributes a cost to every task, we shall assume that each task has a certain utility for every role – and be it just zero. Utility annotations are the basis for the subsequent analysis of incentive alignment, vastly generalizing cost minimization. Note that, in general, it is non-trivial to choose utility functions, especially in competitive situations. However, the O2C process comes with natural candidates for utilities, e.g., postage fees can be looked up from one’s favorite carrier, the cost for gas, maintenance, and personnel for shipping is fairly predictable, and finally there is the profit for selling a good.

A single instance of the O2C process exhibits the phenomenon that *Customer* has no incentive to pay. However, we want to stress that – very much for the same reason – *Shipper* would not have any good reason to perform delivery, once the postage fee is paid. Thus, besides the single instance scenario, we shall consider an unbounded number of repetitions of the process, but only one active process instance at each point in time.⁴ In the repeating variant, the rational reason for the shipper to deliver (and return damaged goods) is expected revenue from future process instances.

One distinguishing feature of the O2C collaboration is that participants do not have to make any joint decisions. Let us illustrate the point with another example. Alice and Bob are co-founders of a company, which is running so smoothly that it suffices when, any day of the week, only one of them is going to work.

Alice suggests that their secretary Mrs. Medina could help them out by rolling a 10-sided die each morning and notifying them about who is going to go to work that day, dependent on whether the outcome is smaller or larger than six. This elaborate process (as shown in Fig. 2), lets Bob work 60% and Alice 40% of the days, respectively. Alice’s reasoning behind it is the observation that Alice is 50% more efficient than Bob when it comes to generating revenue, as indicated by the amount of \$ signs in the process.

In game theoretic terminology, Mrs. Medina is taking the role of a common source of randomness that is independent of the state of the game and does not

⁴ We leave the very interesting situation of interleaved execution of several process instances for future work.

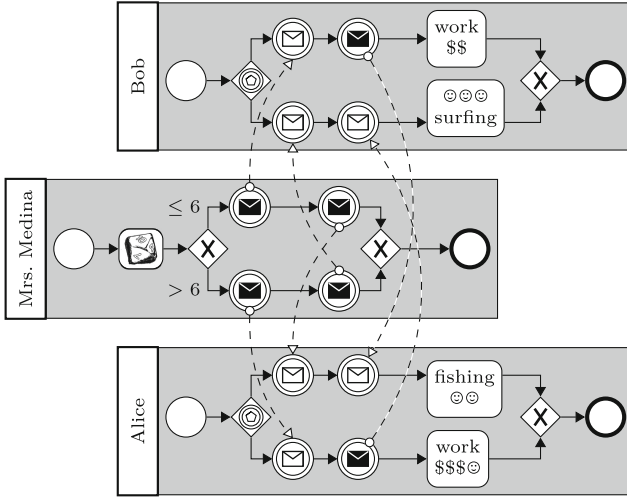


Fig. 2. The *To work or not to work?* collaboration

need to observe the actions of the players. The specific formal notion that we shall use is that of an *autonomous correlation device* [25, Definition 2.1].

Definition 2 (Autonomous correlation device). An autonomous correlation device is a family of pairs $\mathcal{D} = \{(\{M_n^i\}_{i \in N}, d_n)\}_{n \in \mathbb{N}}$ (that is indexed over natural numbers $n \in \mathbb{N}$) each of which consists of

- a family of finite sets of signals M_n^i , (additionally) indexed over players; and
- a function d_n that maps lists of signal vectors $\langle \underline{x}_1, \dots, \underline{x}_{n-1} \rangle \in \prod_{k=1}^{n-1} \mathcal{M}_k$ to probability distributions $d_n \langle \underline{x}_1, \dots, \underline{x}_{n-1} \rangle \in \Delta(\mathcal{M}_n)$ over the Cartesian product $\mathcal{M}_n = \prod_{i=1}^{|N|} M_n^i$ of all signal sets M_n^i .

We shall refer to operators of autonomous correlation devices as *mediators*, which guide the actions of players during the game.

Each correlation device for a game induces an extended game, which proceeds in *stages*. In general, given a game and an autonomous correlation device, the n -th stage begins with the mediator drawing a signal vector $\underline{x}_n \in \mathcal{M}_n = \prod_{i=1}^{|N|} M_n^i$ according to the device distribution $d_n \langle \underline{x}_1, \dots, \underline{x}_{n-1} \rangle$ – e.g., Mrs. Medina rolling the die – and sending the components to the respective players – the sending of messages to Bob and Alice (in one order or the other). Then, each player i chooses an available action a_n^i . This choice can be based on the respective component x_n^i of the signal vector $\underline{x}_n \in \mathcal{M}_n$, information about previous states s_k of the game G , and moves a_k^j of (other) players from the history.⁵ After all players made their choice, we obtain an action profile $a_n = \langle a_n^1, \dots, a_n^{|N|} \rangle$.

⁵ In the present paper, we only consider games of perfect information, which is suitable for business processes in a single organization or which are monitored on a blockchain.

While playing the extended game described above, each player makes observations about the state and the actions of players; the role of the mediator is special insofar as it does not need and is also not expected to observe the run of the game. The “local” observations of each player are the basis of their strategies.

Definition 3 (Observation, strategy, strategy profile). *An observation at stage n by player i is a tuple $h = \langle s_1, \underline{x}_1^i, a_1, \dots, s_{n-1}, \underline{x}_{n-1}^i, a_{n-1}, s_n, \underline{x}_n^i \rangle$ with*

- one state s_k , signal \underline{x}_k^i , and action profile a_k , for each number $k < n$,
- the current state s_n , also denoted by s_h , and
- the current signal \underline{x}_n^i .

The set of all observations is denoted by $H_n^i(\mathcal{D})$. The union $H^i(\mathcal{D}) = \bigcup_{n \in \mathbb{N}} H_n^i(\mathcal{D})$ of observations at any stage is the set of observations of player i . A strategy is a map $\sigma^i: H^i(\mathcal{D}) \rightarrow \Delta(A^i)$ from observations to probability distributions over actions that are available at the current state of histories, i.e., $\sigma_h^i(a^i) = 0$ if $a^i \notin A^i(s_h)$, for all histories $h \in H^i(\mathcal{D})$. A strategy profile is a player-indexed family of strategies $\{\sigma^i\}_{i \in N}$.

Thus, each of the players observes the history of other players, including the possibility of punishing other players for not heeding the advice of the mediator. This is possible since signals might give (indirect) information concerning the (mis-)behavior of players in the past, as remarked by Solan and Vieille [25, p. 370]: by revealing information about proposed actions of previous rounds, players can check for themselves whether some player has ignored some signal of the mediator.

The data of a game, a correlation device, and a strategy profile induce probabilities for finite plays of the game, which in turn determine the expected utility of playing the strategy. Formally, an autonomous correlation device and a strategy profile with strategies for every player yield a probabilistic trajectory of a sequence of “global” states, signal vectors of all players, and complete action profiles, dubbed *history*. The formal details are as follows.

Definition 4 (History and its probability). *A history at stage n is a tuple $h = \langle s_1, \underline{x}_1, a_1, \dots, s_{n-1}, \underline{x}_{n-1}, a_{n-1}, s_n, \underline{x}_n \rangle$ that consists of*

- one state s_k , signal vector \underline{x}_k , and action profile a_k , for each number $k < n$,
- the current state s_n , often denoted by s_h , and
- the current signal vector \underline{x}_n .

The set of all histories at state n is denoted by $H_n(\mathcal{D})$. The union $H(\mathcal{D}) = \bigcup_{n \in \mathbb{N}} H_n(\mathcal{D})$ of histories at arbitrary stages is the set of finite histories. The probability of a finite history $h = \langle s_1, \underline{x}_1, a_1, \dots, s_{n-1}, \underline{x}_{n-1}, a_{n-1}, s_n, \underline{x}_n \rangle$ in the context of a correlation device \mathcal{D} , an initial state s , and a strategy profile σ is defined as follows, by recursion over the length of histories.

$$n = 1: \mathbf{P}_{\mathcal{D}, s, \sigma}(\langle s_1, \underline{x}_1 \rangle) = \begin{cases} 0 & \text{if } s \neq s_1 \\ d_1(\cdot)(\underline{x}_1) & \text{otherwise} \end{cases}$$

$$n > 1: \mathbf{P}_{\mathcal{D},s,\sigma}(\langle \bar{h}, a_{n-1}, s_n, \underline{x}_n \rangle) = \underbrace{p(\bar{h})(a_{n-1})}_{\prod_{i \in N} \sigma_{\langle \bar{h} \rangle}^i(a_{n-1}^i)} q(s_n \mid s_{n-1}, a_{n-1}) \underbrace{p_{d_{n-1}}(\underline{x}_n)}_{d_{n-1}(\underline{x}_1, \dots, \underline{x}_{n-1})(\underline{x}_n)}$$

Again, note that the autonomous correlation device does not “inspect” the states of a history, in the sense that the distributions over signal vectors d_n are *not* parameterized over *states* from the history, but only over previously drawn *signal vectors* – whence the name.

Definition 5 (Mean expected payoff). *The mean expected payoff of player i for stage n is $\bar{\gamma}_n^i(\mathcal{D}, s, \sigma) = \sum_{h \in H_{n+1}(\mathcal{D})} \frac{\mathbf{P}_{\mathcal{D},s,\sigma}(h)}{n} \sum_{k=1}^n u^i(s_k, a_k)$ where $h = \langle s_1, \underline{x}_1, a_1, \dots, a_n, s_{n+1}, \underline{x}_{n+1} \rangle$.*

At this point, we can address the question of what a good strategy profile is and fill in all the details of the idea that an equilibrium is a strategy profile that does not give players any good reason to deviate unilaterally. We shall tip our hats to game theory and use the notation (π^i, σ^{-i}) for the strategy profile which is obtained by “overwriting” the single strategy σ^i of player i with a strategy π^i (which might, but does not have to be different); thus, the expression ‘ (π^i, σ^{-i}) ’ denotes the unique strategy subject to equations $(\pi^i, \sigma^{-i})^i = \pi^i$ and $(\pi^i, \sigma^{-i})^j = \sigma^j$ (for $i \neq j$).

Definition 6 (Autonomous correlated ε -equilibrium). *Given a positive real $\varepsilon > 0$, an autonomous correlated ε -equilibrium is a pair $\langle \mathcal{D}, \sigma^* \rangle$, which consists of an autonomous correlation device \mathcal{D} and a strategy profile σ^* for which there exists a natural number $n_0 \in \mathbb{N}$ such that for any alternative strategy σ^i of any player i , the following inequality holds, for all $n \geq n_0$ and all states $s \in S$.*

$$\bar{\gamma}_n^i(\mathcal{D}, s, \sigma^*) \geq \bar{\gamma}_n^i(\mathcal{D}, s, (\sigma^i, \sigma^{*-i})) - \varepsilon \tag{1}$$

Thus, a strategy is an autonomous correlated ε -equilibrium if the benefits that one might reap in the long run by unilateral deviation from the strategy are negligible as ε can be arbitrarily small. In fact, other players will have ways to punish deviation from the equilibrium [25, § 3.2].

2.2 Petri Nets and Their Operational Semantics

We shall use the definitions concerning Petri nets that have become established in the area of business processes management [2].

Definition 7 (Petri net, marking, and marked Petri net). *A Petri net is a triple $\mathcal{N} = (P, T, F)$ that consists of*

- a finite set of places P ;
- a finite set of transitions T that is disjoint from places, i.e., $T \cap P = \emptyset$; and
- a finite set of arcs $F \subseteq (P \times T) \cup (T \times P)$ (a.k.a. the flow relation).

An input place (resp. output place) of a transition $t \in T$ is a place $p \in P$ s.t. $(p, t) \in F$ (resp. $(t, p) \in F$). The pre-set $\bullet t$ (resp. post-set $t\bullet$) of a transition $t \in T$ is the set of all input places (resp. output places), i.e.,

$$\bullet t = \{p \in P \mid p \text{ is an input place of } t\} \quad t\bullet = \{p \in P \mid p \text{ is an output place of } t\}.$$

A marking of a Petri net \mathcal{N} is a multiset of places m , i.e., a function $m: P \rightarrow \mathbb{N}$ that assigns to each place $p \in P$ a non-negative integer $m(p) \geq 0$. A marked Petri net is a tuple $\mathcal{N} = (P, T, F, m_0)$ whose first three components (P, T, F) are a Petri net and whose last component m_0 is the initial marking, which is a marking of the latter Petri net.

One essential feature of Petri nets is the ability to execute several transitions concurrently – possibly several occurrences of one and the same transition. However, we shall only encounter situations in which a set of transitions fires. To avoid proliferation of terminology, we shall use the general term *step*. We fix a Petri net $\mathcal{N} = (P, T, F)$ for the remainder of the section.

Definition 8 (Step, step transition, reachable marking). A step in the net \mathcal{N} is a set of transitions $\underline{t} \subseteq T$. The transition relation of a step $\underline{t} \subseteq T$ relates a marking m to another marking m' , in symbols $m [\underline{t}] m'$, if the following two conditions are satisfied, for every place $p \in P$.

1. $m(p) \geq |\{t \in \underline{t} \mid p \in \bullet t\}|$
2. $m'(p) = m(p) - |\{t \in \underline{t} \mid p \in \bullet t\}| + |\{t \in \underline{t} \mid p \in t\bullet\}|$

We write $m [\underline{t}] m'$ if $m [\underline{t}] m'$ holds for some step \underline{t} and denote the reflexive transitive closure of the relation $[\underline{t}]$ by $[\underline{t}]^*$. A marking m' is reachable in a marked Petri net $\mathcal{N} = (P, T, F, m_0)$ if $m_0 [\underline{t}]^* m'$ holds, in the net (P, T, F) .

For a transition $t \in T$, we write $m [t] m'$ instead of $m [\{t\}] m'$. Thus the empty step is always fireable, i.e., for each marking m , we have an “idle” step $m [\emptyset] m$.

Recall that a marked Petri net $\mathcal{N} = (P, T, F, m_0)$ is *safe* if all reachable markings m' have at most one token in any place, i.e., if they satisfy $m'(p) \leq 1$, for all $p \in P$. Thus, a marking m corresponds to a set $\hat{m} \subseteq P$ satisfying $p \in \hat{m}$ iff $m(p) > 0$; for convenience, we shall identify a safe marking m with its set of places \hat{m} . The main focus will be on Petri nets that are safe and *extended free choice*, i.e., if the pre-sets of two transitions have a place in common, the pre-sets coincide. Also, recall that the *conflict relation*, denoted by $\#$, relates two transitions if their pre-sets intersect, i.e., $t \# t'$ if $\bullet t \cap \bullet t' \neq \emptyset$, for $t, t' \in T$; for extended free choice nets, the conflict relation is an equivalence relation. We call a marked Petri net an *elementary net system* [23] if all pre-sets and post-sets of transitions are non-empty and every place is input or output to some transition. The latter encompass the following class of Petri nets that is highly relevant to formal methods research of business processes.

Definition 9 (Workflow net (WF-net)). A Petri net $\mathcal{N} = (P, T, F)$ is a Workflow net or WF-net, for short, if

1. *there are unique places $i, o \in P$ such that i is not an output place of any transition and o is not an input place of any transition and*
2. *if we add a new transition t^* and the two arcs $(o, t^*), (t^*, i)$, the resulting directed graph $(P \cup T \cup \{t^*\}, F \cup \{(o, t^*), (t^*, i)\})$ is strongly connected.*

Finally, let us recall the soundness property [1]. A Workflow net is

sound if and only if the following three requirements are satisfied: (1) *option to complete*: for each case it is always still possible to reach the state which just marks place *end*, (2) *proper completion*: if place *end* is marked all other places are empty for a given case, and (3) *no dead transitions*: it should be possible to execute an arbitrary activity by following the appropriate route

where *end* is place *o*, *each case* means every marking reachable from the initial marking $\{i\}$, *state* means marking, *marked* means marked by a reachable marking, *activity* means *transition*, and *following the appropriate route* means after executing the appropriate firing sequence.

3 Incentive Alignment

Soundness of business processes in the sense of Van der Aalst [2] implies termination if transitions are governed by a strongly fair scheduler [1]; indeed, such a scheduler fits the intra-organizational setting. However, as discussed for the O2C process model, unfair scheduling practices could arise in the inter-organizational setting if undesired behavior yields higher profits. We consider incentive alignment to rule out scenarios that lure actors into counterproductive behavior. We even can check whether all activities in a given BPMN model with utility annotations are relevant and profitable.

As BPMN models have established Petri net semantics [6], it suffices to consider the latter for the game theoretic aspects of incentive alignment. As a preparatory step, we extend Petri nets with utility functions as pioneered by von Neumann and Morgenstern [19]. Then we describe two ways to associate a stochastic game to a Petri net with transition-based utilities: the first game retains the state space and the principal design choice concerns transition probabilities; the second game is the restarting version of the first game. Finally, we define incentive alignment in formally based on stochastic games and show that the soundness property for Workflows nets [2] can be “rediscovered” as a special case of incentive alignment; in other words, the original meaning of soundness is conserved, and thus we extend soundness conservatively in our framework for incentive alignment.

3.1 Petri Nets with Utility and Role Annotations

We assume that costs (respectively profits) are incurred (resp. gained) per task and that, in particular, utility functions do not depend on the state. Note that the

game theoretic results do not require this assumption; however, this assumption does not only avoid clutter, but also retains the spirit of the ABC method [12] and is in line with the work of Herbert and Sharp [9].

Definition 10 (Petri net with transition payoffs and roles). For a set of roles \mathcal{R} , a Petri net with transition payoffs and roles is a triple (\mathcal{N}, u, ρ) where

- $\mathcal{N} = (P, T, F, m_0)$ is a marked Petri net with initial marking m_0 ,
- $u: \mathcal{R} \rightarrow T \rightarrow \mathbb{R}$ is a utility function, and
- $\rho: T \rightarrow \mathcal{R}$ is a partial function, assigning at most one role to each transition.

The utility $u^i(\underline{t})$ of a step $\underline{t} \subseteq T$ is the sum of the utilities of its elements, i.e., $u^i(\underline{t}) = \sum_{t \in \underline{t}} u^i(t)$, for each role $i \in \mathcal{R}$.

As a consequence of the definition, the idle step has zero utility. We have included the possibility that some of the transitions are not controlled by any of the roles (of a BPMN model) by using a partial function from transitions to roles; we take a leaf out of the game theorist’s book and attribute the missing role to *nature*.




Fig. 3. Extending Petri nets with role and utility annotations

Figure 3 displays a Petri net on the left. The names of the places p_1, \dots, p_4 will be convenient later. In the same figure on the right, we have added annotations that carry information concerning roles, costs, and profits in the form of lists of role-utility pairs next to transitions. E.g., the transition t_0 is assigned to role a and firing t_0 results in utility -1 for a , i.e., one unit of cost. The first role in each list denotes responsibility for the transition and we have omitted entries with zero utility. We also have colored transitions with the same color as the role assigned to it. If we play the token game for Petri nets as usual, each firing sequence gives cumulative utilities for each one of the roles; each transition gives an immediate reward. These rewards will influence the choice between actions that are performed by roles as made precise in the next subsection.

There are natural translations from BPMN models with payoff annotations for activities to Petri nets with payoffs and roles (relative to any of the established Petri net semantics for models in BPMN [6]). If pools are used, we take one role per pool and each task is assigned to its enclosing pool; for pairs of sending and receiving tasks or events, the sender is responsible for the transition to be taken. The only subtle point concerns the role of nature. When should we blame nature for the data on which choices are based? The answer depends on the application at hand. For instance, let us consider the O2C model of Fig. 1: whether or not

the goods will be damaged during shipment is only partially within the control of the shipper; thus, we shall blame nature for any damage or praise her if everything went well against all odds. In a first approximation, we simply let nature determine whether goods will arrive unscathed.

3.2 Single Process Instances and the Base Game with Fair Conflicts

We now describe how each Petri net with transition payoffs and roles gives rise to a stochastic game, based on two design choices: each role can execute only one (enabled) transition at a time and conflicts are resolved in a probabilistically fair manner. For example, for the net on the right in Fig. 3, we take four states p_0, p_1, p_2, p_3 , one for each reachable marking. The Petri net does not prescribe what should happen if roles a and c both try to fire transitions t_1 and t' simultaneously if the game is in state p_2 . The simplest probabilistically fair solution consists of flipping a coin; depending on the outcome, the game continues in state p_1 or in state p_3 . For the general case, let us fix a safe, extended free-choice net (\mathcal{N}, u, ρ) with payoffs and roles whose initial marking is m_0 where the marked net \mathcal{N} is an elementary net system (e.g., a WF-net).

Definition 11 (The base game with fair conflicts). *Let $\mathcal{X} \subseteq \wp T$ be the partitioning of the set of transitions into equivalence classes of the conflict relation on the set of transitions, i.e., $\mathcal{X} = \{\{t' \in T \mid t' \# t\} \mid t \in T\}$; its members are called conflict sets. Given a safe marking $m \subseteq P$ and a step $\underline{t} \subseteq T$, a maximal m -enabled sub-step is a step \underline{t}' that is enabled at the marking m , is contained in the step \underline{t} , and contains one transition of each conflict set that has a non-empty intersection with the step, i.e., such that all three of $m \upharpoonright \underline{t}'$, $\underline{t}' \subseteq \underline{t}$ and $|\underline{t}'| = |\{X \in \mathcal{X} \mid \underline{t} \cap X \neq \emptyset\}|$ hold. We write $\underline{t}' \sqsubseteq_m \underline{t}$ if the step \underline{t}' is a maximal m -enabled sub-step of the step \underline{t} .*

The base game with fair conflicts $\langle N, S, A, q, u \rangle$ of the net (\mathcal{N}, u, ρ) is defined as follows.

- The set of players $N := \mathcal{R} \cup \{\perp\}$ is the set of roles and nature, $\perp \notin \mathcal{R}$.
- The state space S is the set of reachable markings, i.e., $S = \{m' \mid m_0 \xrightarrow{*} m'\}$.
- The action set of an individual player i is $A^i := \{\emptyset\} \cup \{\{t\} \mid t \in T, \rho(t) = i\}$, which consists of the empty set and possibly singletons of transitions, where $\rho(t) = \perp$ if $\rho(t)$ is not defined. We identify an action profile $a \in A = \prod_{i=1}^{|\mathcal{N}|} A^i$ with the union of its components $a \equiv \bigcup_{i \in N} a^i$.
- In a given state m , the available actions of player i are the enabled transitions, i.e., $A^i(m) = \{\{t\} \in A^i \mid m \upharpoonright \{t\}\}$.
- $q(m' \mid m, \underline{t}) = \sum_{\underline{t}' \sqsubseteq_m \underline{t} \text{ s.t. } m \upharpoonright \underline{t}' = m'} \prod_{X \in \mathcal{X} \text{ s.t. } \underline{t} \cap X \neq \emptyset} \frac{1}{|\underline{t} \cap X|}$
- $u^i(m, \underline{t}) = \sum_{t \in \underline{t}} u^i(t)$ if $i \in \mathcal{R}$ and $u^\perp(m, \underline{t}) = 0$, for all $\underline{t} \subseteq T$, and $m \subseteq P$.

Let us summarize the stochastic game of a given Petri net with transition payoffs and roles. The stochastic game has the same state space as the Petri net, i.e., the set of reachable markings. The available actions for each player at a given marking are the enabled transitions that are assigned to the player, plus

the “idle” step. Each step comes with a state-independent payoff, which sums up the utilities of each single transition, for each player i . In particular, if all players chose to idle, the corresponding action profile is the empty step \emptyset , which gives 0 payoff. The transition probabilities implement the idea that all transitions of an action profile get a fair chance to fire, even if the step contains conflicting transitions. Let us highlight the following two points for a fixed marking and step: (1) given a maximal enabled sub-step, we roll a fair “die” for each conflict set where the “die” has one “side” for each transition in the conflict set that also belongs to the sub-step (unless the “die” has zero sides); (2) there might be several choices of maximal enabled sub-steps that lead to the same marking. In the definition of transition probabilities, the second point is captured by summation over maximal enabled sub-steps of the step and the first point corresponds to a product of probabilities for each outcome of “rolling” one of the “dice”.

We want to emphasize that if additional information about transition probabilities are known, it should be incorporated. In a similar vein, one can adapt the approach of Herbert and Sharp [9], which extends the BPMN language with probability annotations for choices. However, as we are mainly interested in *a priori* analysis, our approach might be preferable since it avoids arbitrary parameter guessing. The most important design choice that we have made concerns the role of nature, which we consider as absolutely neutral; it is not even concerned with progress of the system as it does not benefit from transitions being fired.

Now, let us consider once more the O2C process. If the process reaches the state in which customer’s next step is payment, there is no incentive for paying. Instead, customer can choose to idle, *ad infinitum*. In fact, this strategy yields maximum payoff for the customer. The BPMN-model does not give any means for punishing customer’s payment inertia. However, even earlier there is no incentive for shipper to pick up the goods. Incentives in the single instance scenario can be fixed, e.g., by adding escrow. However, in the present paper, we shall give yet a different perspective: we repeat the process indefinitely.

3.3 Restarting the Game for Multiple Process Instances

The single instance game from Definition 11 has one major drawback. It allows to analyze only a single instance of a business process. We shall now consider a variation of the stochastic game, which addresses the case of multiple instances in the simplest form. The idea is the same as the one for looping versions of Workflow nets that have been considered in the literature, e.g., to relate soundness with liveness [1, Lemma 5.1]: we simply restart the game in the initial state whenever we reach a final marking.

Definition 12 (Restart game). *A safe marking $m \subseteq P$ is final if it does not intersect with any pre-set, i.e., if $m \cap \bullet t = \emptyset$, for all transitions $t \in T$; we write $m \downarrow$ if the marking m is final, and $m \not\downarrow$ if not. Let $\langle N, S, A, q, u \rangle$ be the base game with fair conflicts of the net (N, u, ρ) . The restart game of the net (N, u, ρ) is the game $\langle N, \mathring{S}, \mathring{A}, \mathring{q}, u \rangle$ with*

$$\begin{aligned}
 & - \mathring{S} = S \setminus \{m'' \subseteq P \mid m'' \downarrow\}; \\
 & - \mathring{q}(m' \mid m, \underline{t}) = \begin{cases} q(m' \mid m, \underline{t}) & \text{if } m' \neq m_0 \\ q(m_0 \mid m, \underline{t}) + \sum_{m'' \downarrow} q(m'' \mid m, \underline{t}) & \text{if } m' = m_0 \end{cases}
 \end{aligned}$$

for all $m, m' \in \mathring{S}$; and the available actions restricted to $\mathring{S} \subseteq S$, i.e., $\mathring{A}^i(s) = A^i(s)$, for $s \in \mathring{S}$.




Fig. 4. Restarting process example

For WF-nets, the variation amounts to identifying the final place with the initial place. The passage to the restart game is illustrated in Fig. 4. The restart game of our example is drastically different from the base game. Player c will be better off “cooperating” and never choosing the action t' , but instead idly reaping benefits by letting players a and b do the work. As a consequence, the transition t' will probably never occur since the responsible role has no interest in executing it. Thus, if we assume that the process *may* restart, the net from Fig. 3 is an example where incentives are aligned w.r.t. completion but not with full liveness.

3.4 Incentive Alignment w.r.t. Proper Completion and Full Liveness

We now formalize the idea that participants want to expect benefits from taking part in a collaboration if agents behave rationally – the standard assumption of game theory. The proposed definition of incentive alignment is in principle of qualitative nature, but it hinges on quantitative information, namely the expected utility for each of the business partners of an inter-organizational process.

Let us consider a Petri net with payoffs (\mathcal{N}, u, ρ) , e.g., the Petri net semantics of a BPMN model. Incentive alignment amounts to existence of equilibrium strategies in the associated restart game $\langle \mathcal{N}, \mathring{S}, \mathring{A}, \mathring{q}, u \rangle$ (as per Definition 12) that eventually will lead to positive utility for every participating player. The full details are as follows.

Definition 13 (Incentive alignment w.r.t. completion and full liveness). *Given an autonomous correlation device \mathcal{D} , a correlated strategy profile σ is eventually positive if there exists a natural number $\bar{n} \in \mathbb{N}$ such that, for all larger natural numbers $n > \bar{n}$, the expected payoff of every player is positive, i.e., for all $i \in N$, $\bar{\gamma}_n^i(\mathcal{D}, m_0, \sigma) > 0$. Incentives in the net (\mathcal{N}, u, ρ) are aligned with*

- proper completion *if, for every positive real $\varepsilon > 0$, there exist an autonomous correlation device \mathcal{D} and an eventually positive correlated ε -equilibrium strategy profile σ of the restart game $\langle N, \hat{S}, \hat{A}, \hat{q}, u \rangle$ such that, for every natural number $\bar{n} \in \mathbb{N}$, there exists a history $h \in H_n(\mathcal{D})$ at stage $n > \bar{n}$ with current state $s_h = m_0$ that has non-zero probability, i.e., $\mathbf{P}_{\mathcal{D}, m_0, \sigma}(h) > 0$;*
- full liveness *if, for every positive real $\varepsilon > 0$, there exist an autonomous correlation device \mathcal{D} and an eventually positive correlated ε -equilibrium strategy profile σ of the restart game $\langle N, \hat{S}, \hat{A}, \hat{q}, u \rangle$ such that, for every transition $t \in T$, for every reachable marking m' , and for every natural number $\bar{n} \in \mathbb{N}$, there exists a history $h = \langle m', \underline{x}_1, a_1, \dots, s_{n-1}, \underline{x}_{n-1}, a_{n-1}, s_n, \underline{x}_n \rangle \in H_n(\mathcal{D})$ at stage $n > \bar{n}$ with $t \in a_{n-1}$ and $\mathbf{P}_{\mathcal{D}, m', \sigma}(h) > 0$.*

Both variations of incentive alignment ensure that all participants can expect to gain profits on average, eventually; moreover, something “good” will always be possible in the future where something “good” is either restart of the game (upon completion) or additional occurrences of every transition.

There are several interesting consequences. First, incentive alignment w.r.t. full liveness implies incentive alignment w.r.t. proper completion, for the case of safe, conflict-free elementary net systems where the initial marking is only reachable via the empty transition sequence; this applies in particular to Workflow nets. Next, note that incentive alignment w.r.t. full liveness implies the soundness property for safe, free-choice Workflow nets. The main insight is that correlated equilibria cover a very special case of strongly fair schedulers, not only for the case of a single player. However, we can even obtain a characterization of soundness in terms of incentive alignment w.r.t. full liveness.

Theorem 1 (Characterization of the soundness property). *Let \mathcal{N} be a Workflow net that is safe and extended free-choice; let $(\mathcal{N}, \rho: T \rightarrow \{\Sigma\}, \underline{1})$ be the net with transition payoffs and roles where Σ is a unique role, $\rho: T \rightarrow \{\Sigma\}$ is the unique total role assignment function, and $\underline{1}$ is the constant utility-1 function. The soundness property holds for the Workflow net \mathcal{N} if, and only if, we have incentive alignment w.r.t. full liveness in $(\mathcal{N}, \rho: T \rightarrow \{\Sigma\}, \underline{1})$.*

The full proof can be found in the extended version [8, Appendix A]. However, let us outline the main proof ideas. The first observations is that, w.l.o.g., schedulers that witness soundness of a WF-net can be assumed to be stochastic; in fact, truly random scheduling is strongly fair (with probability 1). Somewhat more detailed, if a WF-net is sound, the scheduler is the only player and scheduling the next best random transition at every point in time yields maximum payoff for the single player. Now, the random choice of a transition at each point in time is the simplest example of an equilibrium strategy (profile); moreover, no matter what the current reachable state of the net, any transition will occur again with non-zero probability, by soundness of the net.

Conversely, incentive alignment w.r.t. strong liveness entails that the unique player – which we might want to think of as the scheduler – will follow a strategy that will eventually fire a transition of the “next instance” of the “process”. In particular, we always will have an occurrence of an *initial* transition by which we

mean a transition that consumes the unique token from the initial marking. After firing an initial transition (of which there will be one by the structure of the net) we are in a state that does not allow us to fire another initial transition. However, strong liveness entails that it has to occur with non-zero probability again if we follow a witnessing equilibrium strategy (profile). Thus, with probability 1, the “current instance” of the “process” will complete such that we will again be able to fire an initial transition.

Finally, the reader may wonder why we consider the restarting game. First, let us emphasize that the restart *games* are merely a means to an end to reason about incentive alignment of BPMN models with suitable utility annotations by use of their execution semantics, i.e., Petri nets with transition payoffs and roles. If these Petri nets do not have any cycles, one could formalize the idea of incentive alignment using finite extensive form games for which correlated equilibria have been studied as well [27]. However, this alternative approach is only natural for BPMN models *without cycles*. In the present paper we have opted for a general approach, which does not impose the rather strong restriction on nets to be acyclic. Notably, while we work with restart *games*, we derive them from *arbitrary* free-choice safe elementary net systems – i.e., without assuming that the input nets are restarting. The restart game is used to check whether incentives are aligned in the original Petri net with transition payoffs and roles.

4 Conclusions and Future Work

We have described a game theoretic perspective on incentive alignment of inter-organizational business processes. It applies to BPMN collaboration models that have annotations for activity-based utilities for all roles. The main theoretical result is that incentive alignment is a conservative extension of the soundness property, which means that we have described a uniform framework that applies the same principles to intra- and inter-organizational business processes. We have illustrated incentive alignment for the example of the order-to-cash process and an additional example that is tailored to illustrate the game theoretic element of mediators.

The natural next step is the implementation of a tool chain that takes a BPMN collaboration model with annotations, transforms it into a Petri net with transition payoffs and roles, which in turn is analyzed concerning incentive alignment, e.g., using algorithms for solving stochastic games [17]. A very challenging venue for future theoretical work is the extension to the analysis of interleaved execution of several instances of a process.

Acknowledgments. We would like to thank the anonymous referees for their detailed comments and suggestions on a previous version of this paper.

References

1. Van der Aalst, W.M.P., Van Hee, K.M., ter Hofstede, A.H.M., et al.: Soundness of workflow nets: classification, decidability, and analysis. *Formal Asp. Comp.* **23**(3), 333–363 (2011)

2. Aalst, W.M.P.: Verification of workflow nets. In: Azéma, P., Balbo, G. (eds.) ICATPN 1997. LNCS, vol. 1248, pp. 407–426. Springer, Heidelberg (1997). https://doi.org/10.1007/3-540-63139-9_48
3. Aumann, R.J.: Subjectivity and correlation in randomized strategies. *J. Math. Econ.* **1**(1), 67–96 (1974)
4. Cachon, G.P., Netessine, S.: Game theory in supply chain analysis. In: Simchi-Levi, D., Wu, S.D., Shen, Z.J. (eds.) *Handbook of Quantitative Supply Chain Analysis*, pp. 13–65. Springer, Heidelberg (2004). https://doi.org/10.1007/978-1-4020-7953-5_2
5. Cartelli, V., Di Modica, G., Manni, D., Tomarchio, O.: A cost-object model for activity based costing simulation of business processes. In: *European Modelling Symposium* (2014)
6. Dijkman, R.M., Dumas, M., Ouyang, C.: Semantics and analysis of business process models in BPMN. *Inf. Softw. Technol.* **50**(12), 1281–1294 (2008)
7. Dumas, M., Rosa, M.L., Mendling, J., Reijers, H.A.: *Fundamentals of Business Process Management*, 2nd edn. Springer, Heidelberg (2018). <https://doi.org/10.1007/978-3-662-56509-4>
8. Heindel, T., Weber, I.: Incentive alignment of business processes: a game theoretic approach (2020). arXiv e-print 2006.06504
9. Herbert, L., Sharp, R.: Using stochastic model checking to provision complex business services. In: *IEEE International Symposium High-Assurance Systems Engineering* (2012)
10. Aumann, R.J.: Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* **55**(1), 1–18 (1987)
11. Jaśkiewicz, A., Nowak, A.S.: Non-zero-sum stochastic games. In: Basar, T., Zaccour, G. (eds.) *Handbook of Dynamic Game Theory*, pp. 1–64. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-27335-8_8-2
12. Kaplan, R., Atkinson, A.: *Advanced Management Accounting*, 3rd edn. Prentice Hall (1998)
13. Kwiatkowska, M., Norman, G., Parker, D., Santos, G.: Automated verification of concurrent stochastic games. In: McIver, A., Horvath, A. (eds.) *QEST 2018*. LNCS, vol. 11024, pp. 223–239. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99154-2_14
14. Kwiatkowska, M., Parker, D., Wiltsche, C.: PRISM-Games 2.0: a tool for multi-objective strategy synthesis for stochastic games. In: Chechik, M., Raskin, J.-F. (eds.) *TACAS 2016*. LNCS, vol. 9636, pp. 560–566. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_35
15. Leyton-Brown, K., Shoham, Y.: *Essentials of game theory: a concise multidisciplinary introduction*. *Synthesis Lect. AI ML* **2**(1), 1–88 (2008)
16. MacDermed, L., Isbell, C.L.: Solving stochastic games. In: *Conference on Neural Information Processing Systems*, pp. 1186–1194 (2009)
17. MacDermed, L., Narayan, K.S., Isbell, C.L., Weiss, L.: Quick polytope approximation of all correlated equilibria in stochastic games. In: *AAAI Conference* (2011)
18. Mendling, J., Weber, I., Van der Aalst, W.M.P., et al.: Blockchains for business process management - challenges and opportunities. *ACM Trans. Manag. Inf. Syst. (TMIS)* **9**(1), 4:1–4:16 (2018)
19. Morgenstern, O., von Neumann, J.: *Theory of games and economic behavior*. Princeton University Press (1953)
20. Narayanan, V., Raman, A.: Aligning incentives in supply chains. *Harvard Bus. Rev.* **82**, 94–102, 149 (2004)

21. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. MIT Press, Cambridge (1994)
22. Prasad, H.L., LA, P., Bhatnagar, S.: Two-timescale algorithms for learning Nash equilibria in general-sum stochastic games. In: *International Conference on Autonomous Agents and Multiagent Systems* (2015)
23. Rozenberg, G., Engelfriet, J.: Elementary net systems. In: Reisig, W., Rozenberg, G. (eds.) *ACPN 1996*. LNCS, vol. 1491, pp. 12–121. Springer, Heidelberg (1998). https://doi.org/10.1007/3-540-65306-6_14
24. Shapley, L.S.: Stochastic games. *Proc. Natl. Acad. Sci.* **39**(10), 1095–1100 (1953)
25. Solan, E., Vieille, N.: Correlated equilibrium in stochastic games. *Games Econ. Behav.* **38**(2), 362–399 (2002)
26. Solan, E., Vieille, N.: Stochastic games. *Proc. Natl. Acad. Sci.* **112**(45), 13743–13746 (2015)
27. von Stengel, B., Forges, F.: Extensive-form correlated equilibrium: definition and computational complexity. *Math. Oper. Res.* **33**, 1002–1022 (2008)
28. Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., Mendling, J.: Untrusted business process monitoring and execution using blockchain. In: La Rosa, M., Loos, P., Pastor, O. (eds.) *BPM 2016*. LNCS, vol. 9850, pp. 329–347. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45348-4_19
29. Weske, M.: *Business Process Management - Concepts, Languages, Architectures*, 3rd edn. Springer, Heidelberg (2019)
30. Xu, X., Weber, I., Staples, M.: *Architecture for Blockchain Applications*. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-03035-3>