

A theory of agreements and protection

Massimo Bartoletti* Tiziana Cimoli*

Dipartimento di Matematica e Informatica, Università degli Studi di Cagliari, Italy

Roberto Zunino

DISI, Università di Trento and COSBI, Italy

Abstract

We present a theory of contracts. Contracts are interacting processes with an explicit notion of obligations and objectives. We model processes and their obligations as event structures. We define a general notion of *agreement*, by interpreting contracts as multi-player concurrent games. A participant agrees on a contract if she has a strategy to reach her objectives (or make another participant chargeable for a violation), whatever the moves of her adversaries. We then tackle the problem of *protection*. A participant is protected by a contract when she has a strategy to defend herself in all possible contexts, even in those where she has not reached an agreement. We show that, in a relevant class of contracts, agreements and protection mutually exclude each other. We then propose a novel formalism for modelling contractual obligations: event structures with circular causality. Using this model, we show how to construct contracts which guarantee both agreements and protection.

1 Introduction

The lack of precise guarantees about the reliability and security of cloud services is a main deterrent for industries wishing to move their applications and business to the cloud [1]. A key problem is how to drive safe and fair interactions among distributed participants which are possibly mutually distrusted, and have possibly conflicting individual goals. In addition to the well-known difficulties of distributed software systems (distribution, concurrency, heterogeneity, mobility, *etc.*), cloud components and infrastructures are often under the governance of different providers, possibly competing among each other. Analysis and verification techniques can be applied only on the software components under one's control, while no assumptions can be made about the components made available by other providers. Therefore, standard compositional techniques have to

*Work partially supported by Aut. Region of Sardinia under grants L.R.7/2007 CRP2-120 (TESLA), CRP-17285 (TRICS), P.I.A. 2010 Project "Social Glue", and by MIUR PRIN 2010-11 project "Security Horizons".

be adapted to cope with the situation where providers fail to keep the promises made, or even choose not to.

We envision a *contract-oriented computing* paradigm [3], where reliable interactions are driven by contracts which formalise Service-Level Agreements. Contracts specify the behavior of a software component, from the point of view of the interactions it may participate in, and the goals it tries to reach. Differently from most of the approaches based on behavioural types [15], which use contracts only in the “matchmaking” phase, a contract-oriented component is not supposed to be *honest*, in that it may not keep the promises made.

In a contract-oriented application, participants advertise their contracts to some *contract brokers*, which are the contract-oriented analogous of service repositories in the Web Service paradigm. Then, participants wait until the contract broker finds an *agreement* among the contracts in its hands. When this happens, a session is created among the participants involved in the contract, so that they can interact. Agreement is a property of contracts which guarantees that each honest participant may reach her objectives, provided that the other participants cooperate honestly. Moreover, if an honest participant does not reach her goals, then some other participant can be blamed. An external judge may then inspect the contract and the status of the session. In case a violation is found, the judge will eventually provide the prescribed compensations/punishments.

The underlying assumption of this view is that participants trust in the contract broker. In a context populated by attackers, it may happen that a dishonest contract broker creates a fraudulent session, making participants interact in the absence of an agreement. In this way, the contract broker may allow an accomplice to swindle an unaware participant. Note that the accomplice may perform his scam while adhering to his contract, and so he will not be liable for violations.

A crucial problem is how to devise contracts which protect participants from malicious contract brokers, while at the same time allowing honest brokers to find agreements. A good contract should allow a participant to reach her goals in contexts where the other participants are cooperative, and prevent her from performing imprudent actions which could be exploited by malicious participants.

In this paper we propose a foundational model for contracts. We specify the behaviour of participants as event structures [21], a basic semantic model for interactive systems. We then provide a formal definition for the two key notions intuitively introduced above, i.e. *agreement* and *protection*. To do that, we borrow techniques from game theory, by interpreting contracts as multi-player concurrent games. By abstracting away from the concrete details of contract languages, our model is a first step towards a unifying framework for reasoning about contracts, in the same spirit that event structures can be used as an underlying semantics for a variety of concrete models of concurrency.

A first result is that agreement and protection cannot coexist for a broad class of objectives. That is, if we are given the objectives of a set of participants, it is impossible to construct a contract which protects them all, and at the same

time admits an agreement. Roughly, the problem is that, when the offers of the participants mutually depend on their requests, the participant which risks in doing the first step is not protected.

To overcome this negative result, we extend event structures with a new notion of causality. While in classical event structures an action a which causally depends on an action b can only be performed after b , in our extension we also consider a relaxed version of causality, which allows a to happen before b , under the (legally binding) promise that b will be eventually performed.

The main result of the paper is that, using this model for contracts, it is possible for a wide class of objectives to construct a set of contracts which protect their participants and still admit an agreement.

2 Contracts

At an abstract level, contracts are concurrent systems, enriched with a notion of *obligation* (what I must do in a given state) and *objective* (what I expect to obtain in a given state). Event structures (ES) are one of the classical models for concurrency [21]. Notwithstanding the variety of formalisations, ES are at least equipped with an *enabling* relation modelling causality (usually written \vdash), and another relation modelling non-determinism (usually written $\#$). ES can provide a basic semantic model for contractual clauses, by interpreting the enabling $\{b\} \vdash a$ as: “I am obliged to do a *after* you have done b ”.

2.1 Event structures

Event structures [21] have a deep theory. Here we only report some basic definitions, which are needed in our technical development. We assume an enumerable universe of *events*, ranged over by a, b, e, \dots . For a set of events X , the predicate $CF(X)$ is true iff X is *conflict-free*, i.e. $CF(X) \triangleq (\forall e, e' \in X : \neg(e\#e'))$.

Definition 2.1 (Event structure [21]). *An event structure \mathcal{E} is a triple $\langle E, \#, \vdash \rangle$, where (1) E is a set of events, (2) $\# \subseteq E \times E$ is an irreflexive and symmetric conflict relation (3) $\vdash \subseteq \{X \subseteq_{\text{fin}} E \mid CF(X)\} \times E$ is the enabling relation, which is saturated, i.e. $\forall X \subseteq Y \subseteq_{\text{fin}} E. X \vdash e \wedge CF(Y) \implies Y \vdash e$.*

An ES is *finite* when E is finite; it is *conflict-free* when the conflict relation is empty. We shall often use the following shorthands: $X \vdash Y$ for $\forall e \in Y. X \vdash e$, $a \vdash b$ for $\{a\} \vdash b$, and $\vdash e$ for $\emptyset \vdash e$.

Definition 2.2 (Persistent conflict). *An event $e \in E$ is persistently conflictable in \mathcal{E} iff the set $\{\bar{e} \in E \mid e\#\bar{e}\}$ is infinite. A set $X \subseteq E$ is persistently conflictable iff some $e \in X$ is persistently conflictable.*

For a sequence $\sigma = \langle e_0 e_1 \dots \rangle$ in E (possibly infinite), we write $\bar{\sigma}$ for the set of elements in σ ; we write σ_i for the subsequence $\langle e_0 \dots e_{i-1} \rangle$. If $\sigma = \langle e_0 \dots e_n \rangle$, we write σe for the sequence $\langle e_0 \dots e_n e \rangle$. The empty sequence is denoted by ε .

For a set S , we denote with S^* the set of finite sequences over S , and with S^ω the set of finite and infinite sequences over S .

A configuration C is a “snapshot” of the behaviour of the system modeled by an ES, where for each event $e \in C$ it is possible to find a finite justification, *i.e.* a sequence of events containing all the causes of e .

Definition 2.3 (Configuration [21]). For an ES $\mathcal{E} = \langle E, \#, \vdash \rangle$, a set $C \subseteq E$ is a configuration of \mathcal{E} iff $CF(C)$, and

$$\forall e \in C. \exists \sigma = \langle e_0 \cdots e_n \rangle. e \in \bar{\sigma} \subseteq C \wedge \forall i \leq n. \bar{\sigma}_i \vdash e_i$$

The set of all configurations of \mathcal{E} is denoted by $\mathcal{F}_\mathcal{E}$.

Definition 2.4 (LTS of an ES). For an ES \mathcal{E} , the labelled transition system $LTS_\mathcal{E} = \langle \wp_{fin}(E), E, \rightarrow_\mathcal{E} \rangle$ is defined as follows:

$$C \xrightarrow{e}_\mathcal{E} C \cup \{e\} \quad \text{iff } C \vdash e, e \notin C \text{ and } CF(C \cup \{e\})$$

Definition 2.5. For two ES $\mathcal{E}, \mathcal{E}'$, we define $\mathcal{E} \sqcup \mathcal{E}'$ as the pointwise union of $\mathcal{E}, \mathcal{E}'$.

2.2 An event-based model of contracts

A contract (Def. 2.6) specifies the obligations and the objectives of a set of participants. The atomic entities of a contract are the *events*, which are uniquely associated to participants through a labelling π . Obligations are modelled as an event structure. Intuitively, an enabling $X \vdash e$ models the fact that, if all the events in X have happened, then e is an obligation for $\pi(e)$. Such obligation may be discharged only by performing e , or any event in conflict with e . For instance, consider an internal choice between two events a and b . This is modelled by an ES with enablings $\vdash a, \vdash b$ and conflict $a \# b$. After the choice (say, of a), the obligation b is discharged. Objectives are modelled as a function Φ , which associates each participant A and each trace of events σ to a *payoff* $\Phi A \sigma$. We assume a rather coarse notion of payoffs: we only have three possible outcomes which represent, respectively, success (1), failure (-1), and tie (0).

Definition 2.6 (Contract). A contract \mathcal{C} is a 4-tuple $\langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$, where:

- $\mathcal{E} = \langle E, \#, \vdash \rangle$ is an event structure;
- \mathcal{A} is a set of participants (ranged over by A, B, \dots);
- $\pi : E \rightarrow \mathcal{A}$ associates each event with a participant;
- $\Phi : \mathcal{A} \rightarrow E^\omega \rightarrow \{-1, 0, 1\}$ associates each participant and trace with a payoff.

Hereafter, we shall assume that contracts respect two basic requirements. For all $X \vdash e$ in \mathcal{E} , we ask that (i) $\Phi(\pi(e)) \neq \perp$, and (ii) e is not persistently conflictable in \mathcal{E} . Notice that Φ is a partial function (from \mathcal{A} to functions),

hence a contract does not need to define payoffs for all the participants in \mathcal{A} (typically, when A advertises her contract, she will not speculate about the objectives of B). Constraint (i) asks that if a contract defines some obligations for A, then A must also declare in \mathcal{C} her payoffs. Constraint (ii) rules out those ill-formed contracts where some obligations can be persistently discharged.

Example 2.7. *Suppose there are two kids who want to play together. Alice has a toy airplane, while Bob has a bike. Both kids are willing to share their toys, but they do not trust each other. Thus, before starting to play they advertise the following contracts. Alice will lend her airplane only after Bob has allowed her ride his bike. Bob will lend his bike without conditions. We model the events “Alice lends her airplane” and “Bob lends his bike” as a and b , respectively. The obligations of Alice and Bob are modelled by the following event structures:*

$$\mathcal{E}_A : \{b\} \vdash a \qquad \mathcal{E}_B : \emptyset \vdash b$$

The objectives of the two kids are modelled by the functions Φ_A (which establishes Alice’s payoff) and Φ_B (for Bob). Alice has a positive payoff in those traces where b has been performed, while she has a negative payoff when she performs a while not obtaining b in return. The payoffs of Bob are dual. Formally:

$$\Phi_A A = \lambda\sigma. \begin{cases} 1 & \text{if } b \in \bar{\sigma} \\ 0 & \text{if } a, b \notin \bar{\sigma} \\ -1 & \text{otherwise} \end{cases} \qquad \Phi_B B = \lambda\sigma. \begin{cases} 1 & \text{if } a \in \bar{\sigma} \\ 0 & \text{if } b, a \notin \bar{\sigma} \\ -1 & \text{otherwise} \end{cases}$$

Summing up, the contracts of Alice and Bob are $\mathcal{C}_A = \langle \mathcal{E}_A, \mathcal{A}, \pi, \Phi_A \rangle$ and $\mathcal{C}_B = \langle \mathcal{E}_B, \mathcal{A}, \pi, \Phi_B \rangle$, respectively, where $\mathcal{A} = \{A, B\}$, $\pi(a) = A$, and $\pi(b) = B$. \square

Observe that the definition of payoff functions in Def. 2.6 is quite liberal. Indeed, it also allows for uncomputable functions, which are of little use in doing anything with a contract. One may then be interested in considering relevant subclasses of payoff functions, in the same spirit of the rich classification of winning conditions in game theory [9].

Assume that participant A has a sequence $\langle O^0 O^1 \dots \rangle$ of sets of events, and a sequence $\langle R^0 R^1 \dots \rangle$ of the same cardinality. The sets O^i are called *offers*, while R^i are the *requests*. A function Φ is an *Offer-Request (O-R) payoff* for A if, whenever A performs in σ some offer O^i (in whatever order), then σ also contains the corresponding request R^i . For instance, the payoff functions Φ_A and Φ_B in Ex. 2.7 are O-R payoffs for A and B. The offers and the requests of A and B are, respectively $O_A^0 = \{a\} = R_B^0$ and, dually, $O_B^0 = \{b\} = R_A^0$.

Definition 2.8 (Offer-Request payoff). *Let $\pi : E \rightarrow \mathcal{A}$. We say that Φ is a Offer-Request payoff for A iff there exist (possibly infinite) sequences $(O^i)_i$, $(R^i)_i$ of equal cardinality such that for all i , $O^i \subseteq \pi^{-1}(A)$, $\emptyset \neq R^i \subseteq E \setminus \pi^{-1}(A)$, and*

$$\Phi A = \lambda\sigma. \begin{cases} 1 & \text{if } (\exists i. R^i \subseteq \bar{\sigma}) \wedge (\forall j. O^j \subseteq \bar{\sigma} \implies R^j \subseteq \bar{\sigma}) \\ 0 & \text{if } (\forall i. R^i \not\subseteq \bar{\sigma} \wedge O^i \not\subseteq \bar{\sigma}) \\ -1 & \text{otherwise} \end{cases}$$

A contract $\mathcal{C} = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ has O-R payoffs iff Φ is a O-R payoff for all $A \in \mathcal{A}$. If, additionally, all the sets O^i (resp. R^i) are finite for all $A \in \mathcal{A}$, we say that \mathcal{C} has finite offers (resp. finite requests). If Φ has a finite number of finite offers-request, then Φ is finite.

Example 2.9. In [8] contracts are modelled in a variant of CCS which includes prefixing, internal/external choice, and recursion. Consider e.g. a server A which repeatedly offers to her clients a choice between two actions a and b . The client B internally chooses one of his (co-)actions \bar{a} and \bar{b} . This is modelled in [8] as:

$$c_A = \text{rec } X. (a.X + b.X) \quad c_B = \text{rec } Y. (\bar{a}.Y \oplus \bar{b}.Y)$$

In our theory we model c_A and c_B as the contracts \mathcal{C}_A and \mathcal{C}_B , defined below. For all $i \geq 0$, let a_i, b_i be events of A , and let \bar{a}_i, \bar{b}_i be events of B . The event structures of A and B have the following enablings and conflicts, for all $i \geq 0$:

$$\begin{aligned} \mathcal{E}_A : & \bar{a}_i \vdash a_i, \bar{b}_i \vdash b_i, a_i \# b_i \\ \mathcal{E}_B : & \vdash \bar{a}_0, \vdash \bar{b}_0, a_i \vdash \bar{a}_{i+1}, a_i \vdash \bar{b}_{i+1}, b_i \vdash \bar{a}_{i+1}, b_i \vdash \bar{b}_{i+1}, \bar{a}_i \# \bar{b}_i \end{aligned}$$

The payoff of a A is positive in a play σ if A has no obligations; similarly for B .

$$\Phi_P P = \lambda \sigma. \begin{cases} 1 & \text{if } \nexists e \in \pi^{-1}(P). \bar{\sigma} \xrightarrow{e}_{\mathcal{E}_P} \\ -1 & \text{otherwise} \end{cases} \quad \square$$

Given two contracts $\mathcal{C}, \mathcal{C}'$, we denote with $\mathcal{C} \mid \mathcal{C}'$ their composition. If \mathcal{C}' is the contract written by an adversary of \mathcal{C} , then a naïve composition of the two contracts could easily lead to an attack, e.g. when Mallory's contract says that Alice is obliged to give him her airplane. To prevent from such kinds of attacks, contract composition is a partial operation. We do *not* compose contracts which assign payoffs to the same participant, neither those which disagree on the association between events and participants.

Definition 2.10 (Composition of compatible contracts). Two contracts $\mathcal{C} = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ and $\mathcal{C}' = \langle \mathcal{E}', \mathcal{A}', \pi', \Phi' \rangle$ are compatible whenever:

$$\forall e \in \mathcal{E} \cap \mathcal{E}'. e = e' \implies \pi(e) = \pi'(e) \quad (1)$$

$$\forall A \in \mathcal{A} \cup \mathcal{A}'. \Phi(A) = \perp \vee \Phi'(A) = \perp \quad (2)$$

If $\mathcal{C}, \mathcal{C}'$ are compatible, we define their composition as:

$$\mathcal{C} \mid \mathcal{C}' = \langle \mathcal{E} \sqcup \mathcal{E}', \mathcal{A} \cup \mathcal{A}', \pi \sqcup \pi', \Phi \sqcup \Phi' \rangle$$

Two contracts which both assign obligations to A are not compatible.

Lemma 2.11. If $\mathcal{C} = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ and $\mathcal{C}' = \langle \mathcal{E}', \mathcal{A}', \pi', \Phi' \rangle$ are compatible, then:

$$X \vdash e \in \mathcal{E} \wedge X' \vdash e' \in \mathcal{E}' \implies \pi(e) \neq \pi'(e') \wedge e \neq e'$$

Example 2.12. The contracts \mathcal{C}_A and \mathcal{C}_B in Ex. 2.7 are compatible, and their composition is the contract $\mathcal{C} = \mathcal{C}_A \mid \mathcal{C}_B = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ defined as follows:

$$\begin{array}{l} \mathcal{E} : \{b\} \vdash a, \emptyset \vdash b \\ \mathcal{A} : \{A, B\} \\ \pi : \pi(a) = A, \pi(b) = B \end{array} \quad \Phi P = \begin{cases} \Phi_A A & \text{if } P = A \\ \Phi_B B & \text{if } P = B \end{cases} \quad \square$$

2.3 Agreements

A crucial notion on contracts is that of *agreement*. Intuitively, when Alice agrees on a contract \mathcal{C} , then she can safely initiate an interaction with the other participants, and be guaranteed that the interaction will not “go wrong” — even in the presence of attackers. This does not mean that Alice will always succeed in all interactions: in case Bob is dishonest, we do not assume that an external authority (e.g. Bob’s mother) will lend the bike to Alice. We intend that Alice agrees on a contract where, in all the interactions where she does not succeed, then some other participant must be found dishonest. That is, we consider Alice satisfied if she can blame another participant. In real-world applications, a judge may provide compensations to Alice, or impose a punishment to the participant who has violated the contract. Here, we shall not explicitly model the judge, and we shall focus instead on how to formalise the agreement property.

We interpret a contract $\mathcal{C} = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ as a nonzero-sum concurrent multi-player game. The game involves the players in \mathcal{A} concurrently performing events in order to reach the objectives defined by Φ . A *play* of \mathcal{C} is a (finite or infinite) sequence of events of \mathcal{E} . We postulate that the permitted moves after a (finite) sequence of steps σ are exactly the events enabled by \mathcal{E} in $\bar{\sigma}$, i.e. e is permitted in σ iff $\bar{\sigma} \xrightarrow{e} \mathcal{E}$. A *strategy* Σ for A is a function which associates to each finite play σ a set of events of A (possibly empty), such that if $e \in \Sigma(\sigma)$ then σe is still a play. A play $\sigma = \langle e_0 e_1 \dots \rangle$ *conforms* to a strategy Σ for A if, for all $i \geq 0$, if $e_i \in \pi^{-1}(A)$, then $e_i \in \Sigma(\sigma_i)$.

As usual in concurrency, we shall only consider those *fair* plays where an event permanently enabled is eventually performed. Indeed, contracts would make little sense in the presence of unfair plays, because an honest participant willing to perform a promised action could be perpetually prevented (by an unfair scheduler) from keeping her promise.

Definition 2.13 (Fair play). A play $\sigma = \langle e_0 e_1 \dots \rangle$ is fair w.r.t. strategy Σ iff:

$$\forall i \leq |\sigma|. (\forall j : i \leq j \leq |\sigma|. e \in \Sigma(\sigma_j)) \implies \exists h \geq i. e_h = e$$

Our notion of agreement takes into account whether participants behave honestly in their plays. Informally, a participant is *innocent* in a play if she always keeps the promises made. An innocent participant has no persistently enabled events, i.e. all her enabled events are either performed or conflicted.

Definition 2.14 (Innocence). We say that A is innocent in σ iff:

$$\forall i \geq 0. \forall e \in \pi^{-1}(A). (\bar{\sigma}_i \xrightarrow{e} \mathcal{E} \implies \exists j \geq i. e_j \# e \vee e_j = e)$$

If A is not innocent in σ , then we say she is culpable.

There always exist strategies which guarantee A to be innocent in every (fair) play. The greatest of such strategies is the *eager strategy*, which prescribes A to do all her enabled events.

Lemma 2.15. *Say a strategy Σ for A is innocent iff A is innocent in all fair plays which conform to Σ . The eager strategy $\Sigma_A^e = \lambda\sigma. \{e \in \pi^{-1}(A) \mid \bar{\sigma} \xrightarrow{e} \varepsilon\}$ is the greatest innocent strategy for A.*

We now define when a participant *wins* in a play. If A is culpable, then she loses. If A is innocent, but some other participant is culpable, then A wins. Otherwise, if all participants are innocent, then A wins if she has a positive payoff in the play. This is formalised as the function \mathcal{W} in Def. 2.16 below.

Definition 2.16 (Winning play). *Define $\mathcal{W} : \mathcal{A} \rightarrow E^\omega \rightarrow \{1, 0, -1\}$ as:*

$$\mathcal{W}A\sigma = \begin{cases} \Phi A\sigma & \text{if all participants are innocent in } \sigma \\ -1 & \text{if A is culpable in } \sigma \\ +1 & \text{otherwise} \end{cases}$$

For a participant A and a play σ , we say that A wins (resp. loses) in σ iff $\mathcal{W}A\sigma > 0$ (resp. $\mathcal{W}A\sigma < 0$).

We can now define when a participant *agrees* on a contract. Intuitively, A is happy to participate in an interaction regulated by contract \mathcal{C} when she has a strategy Σ which allows her to win in all fair plays conform to Σ . More formally, we say that Σ is *winning* (resp. *losing*) for A iff A wins (resp. loses) in every fair play which conforms to Σ .

Definition 2.17 (Agreement). *A participant A agrees on a contract \mathcal{C} if and only if A has a winning strategy in \mathcal{C} . A contract \mathcal{C} admits an agreement whenever all the involved participants agree on \mathcal{C} .*

Example 2.18. *The contract \mathcal{C} of Ex. 2.12 admits an agreement. The winning strategies for A and B are, respectively:*

$$\Sigma_A(\sigma) = \begin{cases} \{a\} & \text{if } b \in \bar{\sigma} \text{ and } a \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases} \quad \Sigma_B(\sigma) = \begin{cases} \{b\} & \text{if } b \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases}$$

For A, the only fair plays conform to Σ_A are ε and $\langle ba \rangle$. B is culpable in ε , while in $\langle ba \rangle$ the payoff of A is positive. For B, the only fair plays conform to Σ_B are $\langle b \rangle$ and $\langle ba \rangle$. A is culpable in $\langle b \rangle$, while in $\langle ba \rangle$ the payoff of B is positive. \square

Example 2.19. *The contracts in Ex. 2.9 above admit an agreement. The winning strategies for A and B are the eager strategies Σ_A^e and Σ_B^e , respectively. \square*

Example 2.20. Note that Σ_A^e is not necessarily winning for A. For instance, consider the contract with $\vdash a, \vdash b, a \# b, \pi^{-1}(A) = \{a, b\}$, and $\Phi A \sigma = 1$ iff $a \in \bar{\sigma}$. We have that $\Sigma_A^e(\varepsilon) = \{a, b\}$, but A is losing in the fair play $\sigma = \langle b \rangle$. However, A agrees on \mathcal{C} , because the strategy $(\lambda \sigma. \text{if } \bar{\sigma} \xrightarrow{a} \text{ then } \{a\} \text{ else } \emptyset)$ is winning for A.

Example 2.21. The union of two winning strategies is not necessarily a winning strategy. For instance, consider the contract with enablings $\vdash a, \vdash b, \{a\} \vdash a', \{b\} \vdash b'$, and conflicts $a \# b', a' \# b$ (all the events are of participant A). Let:

$$\Sigma_a(\sigma) = \begin{cases} \{a\} & \text{if } \bar{\sigma} \xrightarrow{a} \\ \{a'\} & \text{if } \bar{\sigma} \xrightarrow{a'} \\ \emptyset & \text{otherwise} \end{cases} \quad \Sigma_b(\sigma) = \begin{cases} \{b\} & \text{if } \bar{\sigma} \xrightarrow{b} \\ \{b'\} & \text{if } \bar{\sigma} \xrightarrow{b'} \\ \emptyset & \text{otherwise} \end{cases}$$

and let $\Phi A \sigma$ be positive if either $a, a' \in \bar{\sigma}$, or $b, b' \in \bar{\sigma}$. Both Σ_a and Σ_b are winning strategies for A in \mathcal{C} , but their union $\Sigma = \lambda \sigma. \Sigma_a(\sigma) \cup \Sigma_b(\sigma)$ is not. Indeed, $\Sigma(a) = \{a', b\}$, and so $\sigma = \langle ab \rangle$ is a fair play w.r.t. Σ such that $\Phi A \sigma \leq 0$.

We now define the composition \sqcup of a set of strategies. Unlike for the union of winning strategies, their \sqcup -composition is guaranteed to be winning (Lemma 2.23).

Definition 2.22. For a set of strategies \mathcal{S} , we define the strategy $\sqcup \mathcal{S}$ as:

$$(\sqcup \mathcal{S})(\sigma) = \bigcup \{ \Sigma(\sigma) \mid \Sigma \in \mathcal{S} \wedge \sigma \text{ conforms to } \Sigma \}$$

Lemma 2.23. Let $\mathcal{S} = \{\Sigma_1, \dots, \Sigma_n\}$ be a set of strategies. Then:

- (a) A play σ conforms to $\sqcup \mathcal{S}$ iff σ conforms to Σ_i , for some i .
- (b) If all Σ_i are winning for A in \mathcal{C} , then $\sqcup \mathcal{S}$ is a winning strategy for A in \mathcal{C} .

The following lemma gives a necessary condition for reaching an agreement on a contract with O-R payoffs. The ES must have a configuration containing at least a request set, and all the offers are matched by the respective requests.

Lemma 2.24. Let $\mathcal{C} = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ be a contract with O-R payoff $\Phi A = \lambda \sigma. \phi(\bar{\sigma})$ for A. If A agrees on \mathcal{C} , then there exists $C \in \mathcal{F}_{\mathcal{E}}$ such that $\phi(C) > 0$.

The following theorem establishes a sufficient condition for reaching agreements in conflict-free contracts with O-R payoffs. If there exists a configuration C in \mathcal{C} which contains all the requests of A, then A agrees on \mathcal{C} . Since the ES of \mathcal{C} is conflict-free, if the strategy of A prescribes to do all her enabled events in C , then the other participants are obliged to do their events in C . Eventually, either some participant $B \neq A$ is culpable, or a state is reached where the payoff of A is positive.

Theorem 2.25. Let \mathcal{C} be a contract with O-R payoff for A. If \mathcal{E} is conflict-free and $\bigcup_i R_A^i \subseteq C$ for some $C \in \mathcal{F}_{\mathcal{E}}$, then A agrees on \mathcal{C} .

Example 2.26. Note that conflict-freeness is necessary in Theorem 2.25. Consider e.g. the contract \mathcal{C} with O-R payoff for A given by $O_A^0 = \{a\}$ and $R_A^0 = \{b\}$. Assume that $\pi^{-1}(A) = \{a\}$, and $\pi^{-1}(B) = \{b, b'\}$. The ES of \mathcal{C} has enablings $\vdash a, \vdash b, \vdash b'$, and conflict $b\#b'$. The set $\{a, b\}$ is a configuration, but A does not agree on \mathcal{C} . Indeed, either A does no events, or she performs a . In the first case, A is culpable, while in the second one she will have a negative payoff if B does b' . \square

2.4 Protection

In contract-oriented interactions [3], mutually distrusted participants advertise their contracts to a contract broker. The broker composes contracts which admit an agreement, and then establishes a session among the participants involved in such contracts. When a participant agrees on a contract, she is guaranteed that — even in the presence of malicious participants — no interaction driven by the contract will ever go wrong. At worst, if A does not reach her objectives, then some other participant will be found culpable of an infringement.

This model of interaction works fine under the hypothesis that contract brokers are honest, i.e. they never establish a session in the absence of an agreement among all the participants. Suppose Alice is willing to lend her airplane in exchange of Bob's bike. In her contract, she could promise to lend the airplane (unconditionally), and declare that her objective is to obtain the bike. A malicious contract broker could construct an attack by establishing a session between Alice and Mallory, whose contract just says to take the airplane and give nothing in exchange. Mallory is *not* culpable, because her contract declares no obligations, and so Alice loses.

Formally, a contract \mathcal{C}_A *protects* A if, whatever contract \mathcal{C} is composed with \mathcal{C}_A , A has a way to non-lose in the composed contract.

Definition 2.27 (Protection). A contract \mathcal{C}_A protects participant A if and only if, for all contracts \mathcal{C} compatible with \mathcal{C}_A , A has a non-losing strategy in $\mathcal{C}_A \mid \mathcal{C}$.

Notice that if A agrees with \mathcal{C} , then not necessarily \mathcal{C} protects A. For instance, Mallory could join \mathcal{C} with her contract \mathcal{C}_M , and prevent Alice from borrowing Bob's bike in $\mathcal{C} \mid \mathcal{C}_M$. A sufficient (yet hardly realistic) criterion for protection is to declare nonnegative payoffs for all σ . Less trivially, the following example shows a contract with possible negative payoffs which still offers protection.

Example 2.28. The contract \mathcal{C}_B of Ex. 2.7 does not protect Bob. To prove that, consider e.g. the attacker contract $\mathcal{C}' = \langle \mathcal{E}', \mathcal{A}, \pi, \Phi_{\mathcal{C}'} \rangle$, where \mathcal{A} and π are as in Ex. 2.7, while we define \mathcal{E}' with no enablings, and $\Phi_{\mathcal{C}'}$ is immaterial except for being undefined on B (otherwise \mathcal{C}' and \mathcal{C}_B are not compatible). Consider then the contract $\mathcal{C}' \mid \mathcal{C}_B$. There are only two possible strategies for B:

$$\Sigma_B = \lambda\sigma. \emptyset \qquad \Sigma'_B = \lambda\sigma. \begin{cases} \{b\} & \text{if } b \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases}$$

The strategy Σ_B is losing for B, because B is not innocent under Σ_B . The strategy Σ'_B is losing as well, because in the fair play $\sigma = \langle b \rangle$ we have that $\Phi_B \sigma = -1$, but no participant is culpable in σ . By Def. 2.27, B is not protected by \mathcal{C}_B .

On the other hand, the contract \mathcal{C}_A protects Alice. To show that, consider a contract \mathcal{C} compatible with \mathcal{C}_A . Let Σ_A be the following strategy for A:

$$\Sigma_A = \lambda\sigma. \begin{cases} \{a\} & \text{if } b \in \bar{\sigma} \text{ and } a \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases}$$

Let σ be a fair play in $\mathcal{C} \mid \mathcal{C}_A$ conforming to Σ_A . There are two cases:

- $b \in \bar{\sigma}$. Then, since σ is fair, by definition of Σ_A there must exist i such that $a \in \bar{\sigma}_i$, and so A is innocent in σ . Furthermore, we have that $\Phi_A \sigma = 1$.
- $b \notin \bar{\sigma}$. By definition of \mathcal{C}_A , A is not culpable in σ . Also, since $b \notin \bar{\sigma}$ and $a \notin \bar{\sigma}$, then $\Phi_A \sigma = 0$.

In both cases, Σ_A is non-losing for A. Therefore, \mathcal{C}_A protects A. \square

The following theorem establishes sufficient conditions for protection in contracts with O-R payoffs. Essentially, A is protected if, whenever she enables an offer O_A^i , the corresponding request R_A^i has been already satisfied.

Theorem 2.29. *A contract $\mathcal{C}_A = \langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ with O-R payoffs for A protects A if $\forall i, Y. Y \vdash O_A^i \implies R_A^i \subseteq Y$.*

We now consider a relevant subclass of Offer-Request payoffs, where the requests of all participants mutually depend on their offers. An O-R payoff is *circular* when it is not possible to satisfy requests from all participants without each participant doing some offer (item (3)), and each combination of the requests is covered by a set of offers (item (4)). For instance, the payoffs of Alice and Bob in Ex. 2.7 are circular, because their requests (e.g. a and b , respectively) match exactly their offers.

Definition 2.30. *An O-R payoff Φ for participants \mathcal{A} is circular when:*

$$\forall \mathcal{J} : \mathcal{A} \rightarrow \mathbb{N}. \exists \mathcal{L} : \mathcal{A} \rightarrow \mathbb{N}. \bigcup_{A \in \mathcal{A}} O_A^{\mathcal{L}A} \subseteq \bigcup_{A \in \mathcal{A}} R_A^{\mathcal{J}A} \quad (3)$$

$$\forall \mathcal{J} : \mathcal{A} \rightarrow \mathbb{N}. \exists \mathcal{L} : \mathcal{A} \rightarrow \mathbb{N}. \bigcup_{A \in \mathcal{A}} O_A^{\mathcal{L}A} \supseteq \bigcup_{A \in \mathcal{A}} R_A^{\mathcal{J}A} \quad (4)$$

Example 2.31 (Dining retailers [5]). *Around a table, n cutlery retailers are about to have dinner. At the center of the table, there is a large dish of food. Despite the food being delicious, the retailers cannot start eating right now. To do that, and follow the proper etiquette, each retailer needs a complete cutlery set, consisting of n pieces of different kinds. Each of the n retailers owns a distinct set of n piece of cutlery, all of the same kind. The retailers start discussing about trading their cutlery, so that they can finally eat.*

We formalise this scenario as follows. Each retailer A_i initially owns n pieces of kind i . For all $j \neq i$, the event $e_{i,j}$ models A_i giving a piece of cutlery

to retailer A_j . Thus, $\pi^{-1}(A_i) = \{e_{i,j} \mid j \neq i\}$. Retailer A_i offers $n - 1$ pieces of his cutlery of kind i in exchange for $n - 1$ pieces of cutlery of the other kinds.

$$O_i = \{e_{i,j} \mid j \neq i\} \quad R_i = \{e_{j,i} \mid j \neq i\}$$

By Def. 2.30, the payoff Φ_i of each retailer is a finite O-R circular payoff. \square

Agreement and protection can coexist in contracts with *infinite* circular O-R payoffs (see Ex. 2.32). Intuitively, when an infinite offer O_A has to match an infinite request R_B , participants A and B may take turns in doing event in $O_A \cup R_B$. This strategy is winning for both participants (hence they have an agreement), and protection follows because no participant completes her offer before receiving the corresponding request.

Example 2.32. Let $\mathcal{C}_A = \langle \mathcal{E}_A, \mathcal{A}, \pi, \Phi_A \rangle$ and $\mathcal{C}_B = \langle \mathcal{E}_B, \mathcal{A}, \pi, \Phi_B \rangle$ be contracts with circular O-R payoffs (with infinite offers/requests) defined as follows:

$$O_A = \{e_i \mid i \in \mathbb{N}\} = R_B \quad R_A = \{\bar{e}_i \mid i \in \mathbb{N}\} = O_B$$

and let $\mathcal{A} = \{A, B\}$, $\pi(e_i) = A$, $\pi(\bar{e}_i) = B$ for all $i \in \mathbb{N}$. Let the ES \mathcal{E}_A and \mathcal{E}_B be defined by the following enablings (and empty conflicts):

$$\mathcal{E}_A : \{\vdash e_0\} \cup \{\bar{e}_i \vdash e_{i+1} \mid i \geq 0\} \quad \mathcal{E}_B : \{e_i \vdash \bar{e}_i \mid i \geq 0\}$$

The contract $\mathcal{C} = \mathcal{C}_A \mid \mathcal{C}_B$ admits an agreement. We prove separately that A and B agree on \mathcal{C} . Let Σ_A^e be the eager strategy for A. Let σ be a fair play of \mathcal{C} conform to Σ_A^e . We prove that A wins in σ . By Lemma 2.15, the strategy Σ_A^e makes A innocent in σ . There are two subcases. If B is not innocent in σ , then A wins. Otherwise, the play σ must be infinite, i.e. $\bar{\sigma} = \{e_i\}_{i \in \mathbb{N}} \cup \{\bar{e}_i\}_{i \in \mathbb{N}}$. Therefore, $R_A \subseteq \bar{\sigma}$, and so A wins. To prove that B has a winning strategy in \mathcal{C} we proceed similarly, by choosing the eager strategy Σ_B^e for B.

We now show that \mathcal{C}_A protects A. Let \mathcal{C}' be compatible with \mathcal{C}_A . The eager strategy Σ_A^e is non-losing for A. Indeed, in every fair play σ conform to Σ_A^e , if there exists $\bar{e}_i \in R_A \not\subseteq \bar{\sigma}$ then $e_{i+1} \in O_A \not\subseteq \bar{\sigma}$, and so $\Phi_A \sigma \geq 0$. To prove that \mathcal{C}_B protects B, we proceed similarly, by choosing the eager strategy Σ_B^e for B. \square

A remarkable feature of *finite* circular payoffs is that, in each play where all participants win, at some point there exists a participant A which has performed all the offers in O_A^i before having obtained all the requests in R_A^i . Intuitively, the participant A which makes this “first step” is not protected. The proof technique exploited by Lemma 2.33 is similar to that used in [11] to prove that fair exchange is impossible without a trusted third party.

Lemma 2.33. Let \mathcal{C} be a contract with finite circular O-R payoffs. If σ is a winning play for all participants in \mathcal{A} , then there exists a prefix η of σ and a participant $A \in \mathcal{A}$ such that $\Phi_A \eta < 0$.

Our main result in this section is Theorem 2.34 below. It states that if a set of contracts with finite circular O-R payoffs admits an agreement, then some of the participants are not protected, and *vice versa*.

Theorem 2.34. *Let $\mathcal{C}_1, \dots, \mathcal{C}_n$ be contracts with circular finite O-R payoffs for A_1, \dots, A_n , respectively. Then, at most one of the following statements is true:*

- (a) $\mathcal{C}_1 \mid \dots \mid \mathcal{C}_n$ admits an agreement;
- (b) for all $i \in 1..n$, \mathcal{C}_i protects A_i .

3 Reconciling agreement with protection

In the previous section we have shown that agreement and protection cannot coexist in a relevant class of contracts (Theorem 2.34). As made evident by Theorem 2.29, to protect herself A must obtain all her requests R_A^i before doing all her offers O_A^i . If all participants adhere to this principle, agreement is not possible. For instance, Alice and Bob in Ex. 2.7 would be protected by contracts with enablings $a \vdash b$ and $b \vdash a$, but no agreement would be possible because nobody risks doing the first step.

To reconcile agreements with protection, A could relax her contract, i.e. she could do a in change of the *promise* of B to do b . In this case A can safely do the first step, because either B does b , or he will be culpable of a contract violation.

To model this kind of “conditional” enabling, we propose an extension of Winskel’s event structures with a new *circular causality* relation (\Vdash). The enabling $b \Vdash a$ (intuitively, “I will do a if you *promise* to do b ”) together with the other prescription $a \Vdash b$ has a configuration where both a and b have happened, despite of the circular dependencies. We call our extension *ES with circular causality* (CES in short).

In Sect. 3.1 we introduce CES and we state some basic properties. In Sect. 3.2 we reformulate our theory of contracts by using CES in place of ES. Finally, in Sect. 3.3 we show how CES allow for reconciling agreement with protection.

3.1 Event structures with circular causality

Definition 3.1. *An event structure with circular causality is an ES enriched with a (saturated) circular enabling relation $\Vdash \subseteq \{X \subseteq_{fn} E \mid CF(X)\} \times E$.*

We conservatively extend the notion of configuration in [21] to deal with circular causality. Intuitively, for all events e_i in the sequence $\langle e_0 \dots e_n \rangle$, e_i can either be \vdash -enabled by its predecessors, or \Vdash -enabled by the *whole* sequence. Note that if C is a finite configuration, and $\{e_0 \dots e_n\}$ is an enumeration of C which satisfies all the enablings, not necessarily $\{e_0 \dots e_{n-1}\}$ is a configuration as well (see *e.g.*, \mathcal{E}_2 in Fig. 1). To reason compositionally about configurations, Def. 3.2 defines a slightly more general notion of configurations.

In an X -*configuration* C , the set C can contain an event e even in the absence of a justification of e through a standard/circular enabling — provided that e belongs to X . This allows, given an X -configuration, to add/remove any event and obtain an Y -configuration, possibly with $Y \neq X$. We shall say that the events in X have been taken “on credit”, to remark the fact that they may have




Figure 1: CES are denoted as directed hypergraphs, where nodes stand for events. An hyperedge from a set of nodes X to node e denotes an enabling $X \circ e$, where $\circ = \vdash$ if the edge has a single arrow, and $\circ = \Vdash$ if the edge has a double arrow. A conflict $a\#b$ is represented by a dotted line between a and b .

been performed in the absence of a causal justification. Configurations (i.e., \emptyset -configurations) play a crucial role, as they represent sets of events where all the credits have been honoured.

Definition 3.2 (Configuration). Let $\mathcal{E} = (E, \#, \vdash, \Vdash)$ be a CES. A conflict-free sequence $\sigma = \langle e_0 \dots e_n \rangle \in E^*$ without repetitions is an X -trace of \mathcal{E} iff:

$$\forall i \leq n. (e_i \in X \vee \bar{\sigma}_i \vdash e_i \vee \bar{\sigma} \Vdash e_i)$$

For all $C, X \subseteq E$ we say that C is an X -configuration of \mathcal{E} iff $CF(C)$ and:

$$\forall e \in C. \exists \sigma \text{ } X\text{-trace. } e \in \bar{\sigma} \subseteq C$$

The set of all X -configurations of \mathcal{E} is denoted by $\mathcal{F}_{\mathcal{E}}(X)$, or just $\mathcal{F}_{\mathcal{E}}$ when $X = \emptyset$.

Example 3.3. Consider the four CES in Fig. 1.

- (1) \mathcal{E}_1 has enablings $\emptyset \vdash a$, $\emptyset \Vdash b$, and conflict $a\#b$. By Def. 3.2 we have $\emptyset, \{a\}, \{b\} \in \mathcal{F}_{\mathcal{E}_1}$, but $\{a, b\} \notin \mathcal{F}_{\mathcal{E}_1}$.
- (2) \mathcal{E}_2 has enablings $\{a\} \vdash b$ and $\{b\} \Vdash a$. Here $\emptyset, \{a, b\} \in \mathcal{F}_{\mathcal{E}_2}$, $\{b\} \in \mathcal{F}_{\mathcal{E}_2}(\{a\})$ and $\{a\} \in \mathcal{F}_{\mathcal{E}_2}(\{b\})$, while neither $\{a\}$ nor $\{b\}$ belong to $\mathcal{F}_{\mathcal{E}_2}(\emptyset)$.
- (3) \mathcal{E}_3 has enablings $\{a\} \vdash b$, $\{a\} \vdash c$, $\{b\} \Vdash a$, $\{c\} \vdash a$, and conflict $b\#c$. The only non-empty configuration of \mathcal{E}_3 is $\{a, b\}$.
- (4) \mathcal{E}_4 has enablings $\{a, b\} \Vdash c$, $\{a, b\} \Vdash d$, $\{c\} \vdash a$, and $\{d\} \vdash b$. We have that $\{a, b, c, d\} \in \mathcal{F}_{\mathcal{E}_4}$. Note that, were one (or both) of the \Vdash turned into a \vdash , then the only configuration would have been the empty one. \square

According to Winskel's axiom of finite causes, all events in a configuration (except those taken on credit) have a finite justification. Thus, an event cannot be justified through an infinite chain of events, i.e. in the CES with enablings $\{e_{i+1}\} \vdash e_i$ for all $i \geq 0$, the set $\{e_i \mid i \geq 0\}$ is *not* a configuration.

The configurations of CES do still enjoy the finiteness and coherence properties of classical ES, though they are not coincidence-free, which is correct from our point of view because of the presence of circular dependencies. A subset A of a set F is *pairwise compatible* iff $\forall e, e' \in \bigcup A. \exists C \in F. e, e' \in C$.

Theorem 3.4. For all CES \mathcal{E} , and for all $X \subseteq E$, the set $\mathcal{F}_\mathcal{E}(X)$ satisfies:

- (Coherence) If F is a pairwise compatible subset of $\mathcal{F}_\mathcal{E}(X)$, then $\bigcup F \in \mathcal{F}$.
- (Finiteness) $\forall C \in \mathcal{F}. \forall e \in C. \exists C_0 \in \mathcal{F}. e \in C_0 \subseteq_{fin} C$

We define below an operational semantics of CES. This is given in terms of an LTS, the states of which are pairs (C, X) . The first element is the set of events occurred so far; the second element is a set of events taken “on credit”.

Definition 3.5 (LTS of a CES). For a CES \mathcal{E} , we define $LTS_\mathcal{E} = \langle S, E, \rightarrow_\mathcal{E} \rangle$, where $S = \wp_{fin}(E) \times \wp_{fin}(E)$, and $\rightarrow_\mathcal{E}$ is defined by the following rule:

$$\frac{e \notin C \quad CF(C \cup \{e\})}{(C, X) \xrightarrow{e}_\mathcal{E} (C \cup \{e\}, \Delta(C, X, e))}$$

where for all $C, X \subseteq E$ and for all $e \in E$, we define:

$$\Delta(C, X, e) = (X \setminus \{x \in X \mid C \cup \{e\} \Vdash x\}) \cup \begin{cases} \{e\} & \text{if } C \cup \{e\} \not\Vdash e \wedge C \not\vdash e \\ \emptyset & \text{otherwise} \end{cases}$$

The set $\Delta(C, X, e)$ defines how credits change when firing e in a play where the current credits are X , and the events C have already been performed.

The following theorem relates traces of $LTS_\mathcal{E}$ to configurations of \mathcal{E} .

Theorem 3.6. For all CES \mathcal{E} , for all $C, X \subseteq E$:

$$C \in \mathcal{F}(X) \iff \forall D \subseteq_{fin} C. \exists X_0 \subseteq X. \exists C_0. D \subseteq C_0 \subseteq C \wedge (\emptyset, \emptyset) \rightarrow^* (C_0, X_0)$$

3.2 Agreement in CES-based contracts

In this section we conservatively extend the contract theory of Sect. 2, by allowing the component \mathcal{E} of a contract to be a CES.

By Def. 3.5, a conflict-free sequence $\langle e_0 e_1 \dots \rangle$ without repetitions uniquely identifies a trace $(\emptyset, \emptyset) \xrightarrow{e_0}_\mathcal{E} (C_1, X_1) \xrightarrow{e_1}_\mathcal{E} \dots$ in $LTS_\mathcal{E}$. We denote with (C_k^σ, X_k^σ) the state of $LTS_\mathcal{E}$ reached after k steps of the sequence σ . A *play* of a contract \mathcal{C} is a (finite or infinite) sequence σ of events such that $(\emptyset, \emptyset) \xrightarrow{\sigma}_\mathcal{E}$. The notions of *strategy* and *conformance* to a strategy are as in Sect. 2.

The key difference between ES-based and CES-based contracts is the notion of innocence. In the ES-based model, a participant A is culpable in a play σ when some event e of A is enabled in σ . Here, in addition to enabled events, we consider obligations those events which can be done “on credit”, under the guarantee that they will be eventually honoured, whatever events are done later on by the other participants. These events are said *prudent*. The definition of prudent strategies and of innocent participants is mutually coinductive. A participant A is innocent in σ when she has no persistently prudent events. Hence, if the strategy of A tells to do all her prudent events, then in all fair plays these events must either become imprudent, or be fired, or be conflicted.

Formally (although a bit counter-intuitively), fired and conflicted events are imprudent: therefore, \mathbf{A} is innocent when all her prudent events eventually become imprudent.

Definition 3.7 (Prudence). *A strategy Σ for \mathbf{A} is prudent if, for all finite plays σ , for all $e \in \Sigma(\sigma)$ such that σe conforms to Σ , and for all fair plays $\sigma' = \sigma e \eta$ conform to Σ where all $\mathbf{B} \neq \mathbf{A}$ are innocent,*

$$\exists k > |\sigma|. X_k^{\sigma'} \cap \pi^{-1}(\mathbf{A}) \subseteq X_{|\sigma|}^{\sigma'}$$

An event e is prudent in σ if there exists a prudent strategy Σ such that σ conforms to Σ and $e \in \Sigma(\sigma)$.

A participant \mathbf{A} is innocent in $\sigma = \langle e_0 e_1 \dots \rangle$ iff:

$$\forall e \in \pi^{-1}(\mathbf{A}). \forall i \geq 0. \exists j \geq i. e \text{ is imprudent in } \sigma_j$$

Example 3.8. *Recall the CES from Ex. 3.3. In \mathcal{E}_1 , both a and b are prudent in the empty play ε , because they are enabled in \emptyset . In \mathcal{E}_2 , a is prudent in ε , while b is not prudent in ε , yet it is prudent in $\langle a \rangle$. Now consider the CES \mathcal{E}_3 , and assume that $\pi(a) = \mathbf{A}$ and $\pi(b) = \pi(c) = \mathbf{B}$. We have that a is not prudent in ε , because if \mathbf{B} chooses to do c , then the credit a can no longer be honoured. Instead, both b and c are prudent in $\langle a \rangle$.*

Notice that, after Lemma 3.9 below, the new definition of innocence conservatively extends that in Def. 2.14. That is, an event enabled in σ is prudent.

Lemma 3.9. *For all σ and $e \notin \bar{\sigma}$, if $\bar{\sigma} \vdash e$ or $\bar{\sigma} \cup \{e\} \Vdash e$ then e is prudent in σ .*

Lemma 3.10. *Let \mathcal{S} be a finite set of prudent strategies. Then, $\bigsqcup \mathcal{S}$ is prudent.*

Lemma 3.11. *For a contract $\mathcal{C} = \langle \mathcal{E}, \dots \rangle$, where \mathcal{E} is a finite CES, the strategy $\Sigma_{\mathbf{A}}^p = \lambda \sigma. \{e \in \pi^{-1}(\mathbf{A}) \mid e \text{ is prudent in } \sigma\}$ is the greatest prudent strategy for \mathbf{A} .*

Lemma 3.12. *$\Sigma_{\mathbf{A}}^p$ is an innocent strategy for \mathbf{A} .*

For a CES \mathcal{E} and a set of events C , we say that e is *reachable from C* iff there exists η such that $e \in \bar{\eta}$ and $(C, \emptyset) \xrightarrow{\eta} (C', \emptyset)$. Theorem 3.13 states that in conflict-free CES which have only circular enablings, the set of prudent events in σ coincides with the set $\mathcal{R}_{\mathcal{E}}^{\bar{\sigma}}$ of events reachable from $\bar{\sigma}$.

Theorem 3.13. *If \mathcal{E} is conflict-free and \vdash -free, then for all plays σ of $\mathcal{C} = \langle \mathcal{E}, \dots \rangle$:*

$$e \in \mathcal{R}_{\mathcal{E}}^{\bar{\sigma}} \iff e \text{ prudent in } \sigma$$

We now refine the notion of winning strategy given in Def. 2.16. The items are similar to the corresponding items in Def. 2.16, except that the definitions of innocence now takes into account the events performed on credit.

Definition 3.14 (Winning play). Define the function \tilde{W} as follows:

$$\tilde{W}A\sigma = \begin{cases} \Phi A\sigma & \text{A is credit-free and all participants are innocent in } \sigma \\ +1 & \text{if A is innocent, and some } B \neq A \text{ is culpable in } \sigma \\ -1 & \text{otherwise} \end{cases}$$

where we say that A is credit-free in σ iff

$$\forall e \in \pi^{-1}(A). \forall i \geq 0. \exists j \geq i. e \notin X_j^\sigma$$

The notions of winning/losing play/strategy, agreement and protection are the same as in Sect. 2, except that \tilde{W} is now used in place of W .

Lemma 3.15. Let Σ_A be a prudent strategy for A. For all fair plays σ conform to Σ_A , either A is credit-free in σ , or some $B \neq A$ is culpable in σ .

Example 3.16. In Ex. 2.28 we have shown that the contract \mathcal{C}_A protects Alice, while \mathcal{C}_B does not protect Bob. Suppose now to change Bob's contract into a contract \mathcal{C}'_B where Bob relaxes his requirements. The contract \mathcal{C}'_B differs from \mathcal{C}_B only in the event structure \mathcal{E}'_B , which contains exactly one circular enabling: $\{a\} \Vdash b$. Similarly to Ex. 2.18, the contract $\mathcal{C}_A \mid \mathcal{C}'_B$ admits an agreement. To show that, let Σ_A and Σ_B be the following strategies for A and B, respectively:

$$\Sigma_A(\sigma) = \begin{cases} \{a\} & \text{if } b \in \bar{\sigma} \text{ and } a \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases} \quad \Sigma_B(\sigma) = \begin{cases} \{b\} & \text{if } b \notin \bar{\sigma} \\ \emptyset & \text{otherwise} \end{cases}$$

Roughly, the only fair play which conforms to Σ_A and Σ_B where both A and B are innocent is $\sigma = \langle ba \rangle$, which gives rise to the following trace in LTS_ε :

$$(\emptyset, \emptyset) \xrightarrow{b} (\{b\}, \{b\}) \xrightarrow{a} (\{a, b\}, \emptyset)$$

We have that A and B win in σ , because $\tilde{W}A\sigma = 1 = \tilde{W}B\sigma$. Thus, Σ_A and Σ_B are winning strategies for A and B, respectively, and so \mathcal{C} admits an agreement. \square

3.3 Protection in CES-based contracts

In this section we show that CES-based contracts allow for both agreements and protection in contracts with circular finite O-R payoffs. Before presenting the formal results, we give some intuition through our working example.

Example 3.17. Differently from the contract \mathcal{C}_B in Ex. 2.28, the contract \mathcal{C}'_B in Ex. 3.16 protects Bob. Let \mathcal{C}' be a contract compatible with \mathcal{C}'_B . Consider the strategy Σ_B^p for B, as defined in Lemma 3.12. Let ν be a fair play of $\mathcal{C}'_B \mid \mathcal{C}'$ conform to Σ_B^p . By contradiction, assume that B loses in ν . By Lemma 3.12, B is innocent in ν , and so it must be $\Phi B\nu < 0$. By definition, the payoff of B is negative only when $b \in \bar{\nu}$ and $a \notin \bar{\nu}$. Assume that $\nu = \eta b \eta'$. By definition

of $\Sigma_{\mathbf{B}}^p$, the event b was prudent in η , and we have the transition $(\bar{\eta}, X^0) \xrightarrow{b} (\bar{\eta} \cup \{b\}, X^0 \cup \{b\})$. After \mathbf{B} has performed b , its only strategy is the empty one. By Def. 3.7, for all plays $e_0 e_1 \dots$ starting from $(\bar{\eta} \cup \{b\}, X^0 \cup \{b\})$, there exists some $k > 0$ such that $b \notin X^k$. This means that b has been honoured, and the only way to do that is to perform a . Therefore, $a \in \bar{\nu}$ — contradiction. \square

We now construct a CES from an O-R payoff with finite responses. For all clauses (O, R) , the CES contains the enablings $R \Vdash O$. Lemma 3.19 below reveals a key feature of circularity: the CES obtained from a circular O-R payoff has a configuration which comprises all the responses of all participants. Together with Theorem 3.20, this will allow for constructing a contract which admits an agreement. Theorems 3.20 and 3.21 are the CES counterpart of Theorems 2.25 and 2.29 for ES-based contracts, respectively.

Definition 3.18. For an O-R payoff Φ with clauses $(O^i, R^i)_i$ and finite R^i , define $\mathcal{E}(\Phi)$ as the conflict-free CES with (saturated) enablings $\{R^i \Vdash O^i\}_i$.

Lemma 3.19. Let Φ be a finite circular O-R payoff for \mathcal{A} such that $\Phi \mathbf{A} = \lambda \sigma. \phi_{\mathbf{A}} \bar{\sigma}$ for all $\mathbf{A} \in \mathcal{A}$. Then, $\exists C \in \mathcal{F}_{\mathcal{E}(\Phi)}. \forall \mathbf{A} \in \mathcal{A}. \bigcup_i R_{\mathbf{A}}^i \subseteq C$.

Theorem 3.20. Let \mathcal{C} be a contract with O-R payoff for \mathbf{A} . If \mathcal{E} is conflict-free and \vdash -free, and $\bigcup_i R_{\mathbf{A}}^i \subseteq C$ for some $C \in \mathcal{F}_{\mathcal{E}}$, then \mathbf{A} agrees on \mathcal{C} .

Theorem 3.21. For a finite CES \mathcal{E} and an O-R payoffs Φ for \mathbf{A} , the contract $\langle \mathcal{E}, \mathcal{A}, \pi, \Phi \rangle$ protects \mathbf{A} if: $\forall i, Y. (\forall e \in O_{\mathbf{A}}^i. Y \vdash e \vee Y \Vdash e) \implies R_{\mathbf{A}}^i \subseteq Y$.

Theorem 3.22 below states that agreements and protection can coexist in CES-based contracts with circular finite O-R payoffs. Recall that Theorem 2.34 excluded this possibility for ES-based contracts. Condition (5) in Theorem 3.22 is technical, yet it makes the theorem applicable to a broad class of contracts with O-R payoffs (e.g. the dining retailers scenario, see Ex. 3.24). Ex. 3.23 shows that when condition (5) is not satisfied, Theorem 3.22 does not hold in general.

Theorem 3.22. Let Φ_1, \dots, Φ_n be finite circular O-R payoffs for $\mathbf{A}_1, \dots, \mathbf{A}_n$, respectively, and such that, for all $\mathbf{A} \in \{\mathbf{A}_1, \dots, \mathbf{A}_n\}$:

$$\forall P \subseteq \mathbb{N}. \forall j. O_{\mathbf{A}}^j \subseteq \bigcup_{i \in P} O_{\mathbf{A}}^i \implies R_{\mathbf{A}}^j \subseteq \bigcup_{i \in P} R_{\mathbf{A}}^i \quad (5)$$

Then, there exist contracts $\mathcal{C}_i = \langle \mathcal{E}_i, \mathcal{A}, \pi, \Phi_i \rangle$ for $i \in 1..n$ such that:

- (a) $\mathcal{C}_1 \mid \dots \mid \mathcal{C}_n$ admits an agreement;
- (b) for all $i \in 1..n$, \mathcal{C}_i protects \mathbf{A}_i .
- (c) for all plays σ of $\mathcal{C}_1 \mid \dots \mid \mathcal{C}_n$, $\forall e \in \bar{\sigma}. \exists i. e \in O_{\pi(e)}^i$.

Example 3.23. Consider the O-R payoff $\Phi_{\mathbf{A}}$ of participant \mathbf{A} defined by:

$$\begin{array}{lll} O^0 = \{a_0, a_1\} & O^1 = \{a_1, a_2\} & O^2 = \{a_0, a_2\} \\ R^0 = \{b_0\} & R^1 = \{b_1\} & R^2 = \{b_0, b_1\} \end{array}$$

Condition (5) of Theorem 3.22 is satisfied, hence the contract with CES $\mathcal{E}(\Phi_A)$ protects A, and allows A to reach an agreement with other participants whenever the overall payoff satisfies the conditions of the theorem.

Suppose now to change Φ_A , by requiring $R^2 = \{b_2\}$. Notice that such modified payoff no longer satisfies condition (5). Indeed, by choosing $P = \{0, 1\}$ and $j = 2$ we have that $\{a_0, a_2\} = O^2 \subseteq O^0 \cup O^1$, but $\{b_2\} = R^2 \not\subseteq R^0 \cup R^1$. So, Theorem 3.22 does not apply. The CES $\mathcal{E}(\Phi_A)$ contains the enablings $\{b_0\} \Vdash \{a_0, a_1\}$, $\{b_1\} \Vdash \{a_1, a_2\}$, and $\{b_2\} \Vdash \{a_0, a_2\}$. Now A is not protected. Indeed, an attacker B could perform b_0 and b_1 to oblige A to do a_0, a_1, a_2 . A would lose, because to be innocent she has to do all the offers in O^2 , but doing so she is not guaranteed to obtain R^2 . As a matter of facts, there exists no CES which guarantees both agreement and protection for the payoff Φ_A . \square

Example 3.24. Recall the dining retailers scenario from Ex. 2.31. The payoff Φ_i of each retailer is a finite O-R circular payoff, and condition (5) is trivially satisfied. Therefore, Theorem 3.22 allows for constructing contracts which admit an agreement and protects all retailers. The CES of contract \mathcal{C}_i of retailer A_i has enablings $\{e_{j,i} \mid i \neq j\} \Vdash \{e_{i,j} \mid i \neq j\}$. The idea is simple: A_1 offers his pieces of cutlery, in exchange of the commitment of the other retailers to do the same. Since all retailers commit to the analogous contract, we have an agreement. \square

4 Related work and conclusions

We have studied contracts from a foundational perspective. Our formalisation of contracts builds upon a very abstract model of concurrent computations, namely event structures, to provide general notions of agreement and protection. We expect that specific formalisations of agreement, e.g. the one in [8], can be interpreted as instances of our general notion, in the same spirit that event structures can provide semantics to more concrete models of concurrency, e.g. CCS, π -calculus and Petri nets [21].

An abstract model of contracts is fundamental for the development of the contract-oriented paradigm. In addition to the possibility of relating different formalisations of contracts, such an abstract model would also allow for reasoning uniformly about the properties of contract-based systems. For instance, the static/dynamic notions of *honesty* of a process, which in [4] were specific for the contracts of [8], could be generalised to a broader class of contracts.

Aiming at generality, we have almost neglected some relevant issues, e.g. devising efficient decision procedures for agreements. Although in the most general setting (infinite event structures, arbitrary payoff functions) we come up against the problem of undecidability, such kind of results can be obtained by considering suitable subclasses of event structures/payoff functions (e.g. model checking temporal logic on finite representations of infinite event structures, as in [18]).

A heterogeneous ecosystem of formalisms for contracts has appeared in the literature. Citing a few recent approaches, these formalisms include logics [5, 19],

behavioural types [6, 8], Petri nets [20], multi-player games [13], domain-specific languages [16], c-semirings [7], *etc.*

Most of the existing models do not explicitly deal with the circularity issue, which instead has been a main subject of study in this paper. An exception is [5], where circularity is dealt with at a proof-theoretic level. The logic PCL presented in [5] extends propositional intuitionistic logic with a new connective, that weakens the standard implication \rightarrow , somehow similarly to the way our \Vdash weakens the standard enabling \vdash . CES and PCL are strongly related: preliminary results suggest that finite conflict-free CES correspond to Horn PCL formulae.

In [2] some preliminary work on event structures with circular causality is presented. In the simplified model of [2], where event structures are finite and conflict-free, and the goals are O-R payoffs without offers, it is shown how to decide agreements through an encoding of event structures into PCL formulae.

In [14] event structures are extended with a *response* relation. A relation $a \bullet \rightarrow b$ models the fact that, whenever event a is present in a trace, then b must eventually occur after a . This is quite different from a circular enabling $a \Vdash b$, which instead does not impose any ordering between a and b (it suffices that b is honoured somehow). Also, augmenting the number of \Vdash -enablings increases the number of configurations, while adding more response relations reduces it.

Liability issues are the focus of [16, 12]. Given a contract and an execution trace, the problem is to establish evidence about the occurrence of a contract violation, and in particular to assign blame to misbehaving participants. While [16, 12] are not concerned about how an agreement is found (they just consider the contract as already agreed upon), they explore issues not explicitly modelled in our framework. The notion of liability in [16] takes into account time constraints. Extending our contract model with temporal deadlines and, more in general, with quantitative aspects (like e.g. probabilities) seems to be feasible, along the lines of analogous extensions of events structures [17].

Our model adopts a draconian notion of innocence, in that a participant omitting to perform a single due event in a play is considered culpable, regardless of the fact that the other participants could equally be satisfied with that play. Establishing finer-grained notions of causality between a violation and the resulting failure, as done e.g. in [12], seems a plausible extension of our work.

Our notion of winning play (Def. 2.16 and Def. 3.14) is a sort of lexicographic objective, similarly to those in [10]. The secure equilibria introduced in [10] require that a player cannot increase her payoff while decreasing the payoff of the other player. This is stronger than our notion of agreement, where we just require that strategies exist which yield a positive payoff for all players. Indeed, such strategies do not even have to form a Nash equilibrium.

References

- [1] M. Armbrust et al. A view of cloud computing. *Comm. ACM*, 53(4):50–58, 2010.
- [2] M. Bartoletti, T. Cimoli, G. M. Pinna, and R. Zunino. An event-based model for contracts. In *Proc. PLACES*, 2012.
- [3] M. Bartoletti, E. Tuosto, and R. Zunino. Contract-oriented computing in CO₂. *Scientific Annals in Computer Science*, 22(1):5–60, 2012.
- [4] M. Bartoletti, E. Tuosto, and R. Zunino. On the realizability of contracts in dishonest systems. In *Proc. COORDINATION*, 2012.
- [5] M. Bartoletti and R. Zunino. A calculus of contracting processes. In *LICS*, 2010.
- [6] M. Bravetti and G. Zavattaro. Towards a unifying theory for choreography conformance and contract compliance. In *Software Composition*, 2007.
- [7] M. G. Buscemi and U. Montanari. CC-Pi: A constraint-based language for specifying service level agreements. In *Proc. ESOP*, 2007.
- [8] G. Castagna, N. Gesbert, and L. Padovani. A theory of contracts for web services. *ACM Transactions on Programming Languages and Systems*, 31(5), 2009.
- [9] K. Chatterjee and T. A. Henzinger. A survey of stochastic ω -regular games. *J. Comput. Syst. Sci.*, 78(2):394–413, 2012.
- [10] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theor. Comput. Sci.*, 365(1-2):67–82, 2006.
- [11] S. Even and Y. Yacobi. Relations among public key signature systems. Technical Report 175, Computer Science Department, Technion, Haifa, 1980.
- [12] G. Gößler, D. L. Métayer, and J.-B. Raclet. Causality analysis in contract violation. In *Proc. RV*, 2010.
- [13] A. S. Henriksen. *Adversarial Models for Cooperative Interactions*. PhD thesis, Department of Computer Science, University of Copenhagen, 2011.
- [14] T. T. Hildebrandt and R. R. Mukkamala. Declarative event-based workflow as distributed dynamic condition response graphs. In *Proc. PLACES*, 2010.
- [15] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL*, 2008.
- [16] T. Hvitved, F. Klaedtke, and E. Zălinescu. A trace-based model for multiparty contracts. *JLAP*, 81(2):72–98, 2012.

- [17] J.-P. Katoen. *Quantitative and qualitative extensions of event structures*. PhD thesis, University of Twente, 1996.
- [18] W. Penczek. Model-checking for a subclass of event structures. In *Proc. TACAS*, volume 1217 of *Lecture Notes in Computer Science*, pages 145–164. Springer, 1997.
- [19] C. Prisacariu and G. Schneider. A formal language for electronic contracts. In *Proc. FMOODS*, 2007.
- [20] W. M. P. van der Aalst, N. Lohmann, P. Massuthe, C. Stahl, and K. Wolf. Multiparty contracts: Agreeing and implementing interorganizational processes. *Comput. J.*, 53(1):90–106, 2010.
- [21] G. Winskel. Event structures. In *Advances in Petri Nets*, pages 325–392, 1986.

A Basic results on Event Structures

Definition A.1 (Families of configurations [21]). For a set of sets \mathcal{F} we define the following three properties:

Coherence If \mathcal{A} is a pairwise compatible subset of \mathcal{F} , then $\bigcup \mathcal{A} \in \mathcal{F}$.

Finiteness

$$\forall C \in \mathcal{F}. \forall e \in C. \exists C_0 \in \mathcal{F}. e \in C_0 \subseteq_{fin} C$$

Coincidence-freeness

$$\forall C \in \mathcal{F}. \forall e, e' \in C. (e \neq e' \implies (\exists C' \in \mathcal{F}. C' \subseteq C \wedge (e \in C' \iff e' \notin C')))$$

A set \mathcal{F} is a family of configurations iff it satisfies coherence, finiteness and coincidence-freeness, and in that case, \mathcal{F} is called a family of configurations of E when $\bigcup \mathcal{F} = E$.

Theorem A.1 ([21]). For all ES \mathcal{E} , the set $\mathcal{F}_{\mathcal{E}}$ is a family of configurations.

Definition A.2 ([21]). Let \mathcal{F} be a family of configurations of a set E . We define the ES $\hat{\mathcal{E}}(\mathcal{F}) = (E, \#, \vdash)$ as follows:

$$(a) e \# e' \iff \forall C \in \mathcal{F}. e \notin C \vee e' \notin C$$

$$(b) X \vdash e \iff CF(X) \wedge X \text{ is finite} \wedge \exists C \in \mathcal{F}. e \in C \subseteq X \cup \{e\}$$

Theorem A.2 ([21]). For all families of configurations \mathcal{F} , we have $\mathcal{F}_{\hat{\mathcal{E}}(\mathcal{F})} = \mathcal{F}$.

Lemma A.1. For all ES \mathcal{E} , and for all $C \subseteq E$:

$$C \in \mathcal{F}_{\mathcal{E}} \iff \forall D \subseteq_{fin} C. \exists C_0. D \subseteq C_0 \subseteq_{fin} C. \emptyset \xrightarrow{\mathcal{E}}^* C_0$$

Lemma A.2. For all $C, C' \subseteq_{fin} E$, and for all $e \in E$ such that $CF(C' \cup \{e\})$:

$$C \xrightarrow{\mathcal{E}} \wedge C \subseteq C' \wedge e \notin C' \implies C' \xrightarrow{\mathcal{E}}$$

B Proofs for Section 2

B.1 Proof of Lemma 2.24

Assume that A agrees on \mathcal{C} . By Def. 2.17, A has a winning strategy in \mathcal{C} . Let Σ_A be such strategy. By Def. 2.16, A wins in every fair play which conforms to Σ_A . Among all such plays, there must exist at least one where all the participants are innocent (e.g. the play where all $B \neq A$ adopt the eager strategy Σ_B^e). Let then σ be a fair play which conforms to Σ_A , and where all the other participants are innocent. Since A wins in σ , by Def. 2.16 we have $\Phi A \sigma > 0$. To conclude, it suffices to observe that for all plays σ , the set $\bar{\sigma}$ is a configuration of \mathcal{E} .

B.2 Proof of Theorem 2.25

Let \mathcal{C} be a contract with O-R payoffs for A. Let $C \in \mathcal{F}_\mathcal{E}$ be such that $\bigcup_i R_A^i \subseteq C$. We will prove that the eager strategy Σ_A^e is winning for A in \mathcal{C} . Let γ be a fair play of \mathcal{C} which conforms to Σ_A^e .

By contradiction, assume that A is *not* winning in γ . By Lemma 2.15, A is innocent in γ . Thus, by Def. 2.16 it follows that all participants are innocent and $\Phi A \gamma \leq 0$. By Def. 2.8, this means that either there exists some i such that $O_A^i \subseteq \bar{\gamma}$ and $R_A^i \not\subseteq \bar{\gamma}$ (in case A loses), or for all i , $O_A^i \not\subseteq \bar{\gamma}$ and $R_A^i \not\subseteq \bar{\gamma}$ (in case A non-loses). In both cases, there exists at least one i such that $R_A^i \not\subseteq \bar{\gamma}$.

Let i be such that $R_A^i \not\subseteq \bar{\gamma}$, and let e be such that $e \in R_A^i \setminus \bar{\gamma}$. By hypothesis, there exists $C \in \mathcal{F}_\mathcal{E}$ such that $\bigcup_i R_A^i \subseteq C$, hence $e \in C$. Since C is a configuration, and since families of configurations enjoy finiteness (Theorem A.1), there exists $C' \subseteq_{\text{fin}} C$ such that $C' \in \mathcal{F}_\mathcal{E}$ and $e \in C'$. By Lemma A.1, there exists a play σ such that $\emptyset \xrightarrow{\sigma} \bar{\sigma}$, and $e \in \bar{\sigma} = C'$.

We will prove that $\bar{\sigma} \subseteq \bar{\gamma}$ by induction on the length of σ . The base case $\sigma_0 = \varepsilon$ is trivial. For the inductive case, by the induction hypothesis assume that $\bar{\sigma}_i \subseteq \bar{\gamma}$ for $i < |\sigma|$. We have to prove that $\bar{\sigma}_{i+1} = \bar{\sigma}_i \cup \{e_i\} \subseteq \bar{\gamma}$. By the induction hypothesis, it is enough to prove that $e_i \in \bar{\gamma}$.

By contradiction, assume that $e_i \notin \bar{\gamma}$. Let γ_k be the shortest prefix of γ such that $\bar{\sigma}_i \subseteq \bar{\gamma}_k$. Since $\bar{\sigma}_i \vdash e_i$, by Lemma A.2 it follows that $\bar{\gamma}_h \xrightarrow{e} \mathcal{E}$ for all $h \geq k$. Since all participants are innocent in γ , and since \mathcal{E} is conflict-free, by Def. 2.14 there exists $j > k$ such that the j -th event of γ is e_i — contradiction.

Summing up, we have proved that $\bigcup_i R_A^i \subseteq \bar{\gamma}$ for all fair plays γ . Therefore, A has a winning strategy (Σ_A^p) in \mathcal{C} , and so we conclude that A agrees on \mathcal{C} . \square

B.3 Proof of Theorem 2.29

Let \mathcal{C}_A be a contract with O-R payoffs for A such that:

$$\forall i, Y. Y \vdash O_A^i \implies R_A^i \subseteq Y \quad (6)$$

Let \mathcal{C} be a contract compatible with \mathcal{C}_A . We will prove that Σ_A^e is a non-losing strategy for A in $\mathcal{C}_A \mid \mathcal{C}$. Let σ be a fair play of $\mathcal{C}_A \mid \mathcal{C}$ conforming to Σ_A^e .

By contradiction, assume that A loses in σ , i.e. by Def. 2.16 and by Def. 2.8:

$$\exists i. O_A^i \subseteq \bar{\sigma} \wedge R_A^i \not\subseteq \bar{\sigma} \quad (7)$$

Since $O_A^i \subseteq \bar{\sigma}$, then for all $e \in O_A^i$ it must be $\bar{\sigma} \vdash e$. Since $\bar{\sigma} \vdash O_A^i$, by (6) it follows that $R_A^i \subseteq \bar{\sigma}$ — which contradicts (7). \square

B.4 Proof of Lemma 2.33

Since σ is a winning play for A, by Def. 2.16 it must be $\Phi A \sigma > 0$ for all A. Let ηe be the shortest prefix of σ such that all the participants have strictly positive payoffs (of course, a winning play must have strictly positive length). We shall prove that $\Phi B \eta < 0$, for some participant B.

By contradiction, assume that there exists no $A \in \mathcal{A}$ such that $\Phi A \eta < 0$. Since not all participants are winning in η , there should then exist some participant B such that $\Phi B \eta = 0$, for all j , $R_B^j \not\subseteq \bar{\eta}$. Since ηe is winning for \mathcal{A} (and in particular for B), let i_B be such that $e \in R_B^{i_B} \subseteq \bar{\eta} \cup \{e\}$, and for all $A \neq B$ let i_A be such that $R_A^{i_A} \subseteq \bar{\eta} \cup \{e\}$. Since Φ is circular, by Def. 2.30 there exists a function $\mathcal{J} : \mathcal{A} \rightarrow \mathbb{N}$ such that $\bar{\eta} \cup \{e\} \supseteq \bigcup_{A \in \mathcal{A}} R_A^{i_A} \supseteq \bigcup_{A \in \mathcal{A}} O_A^{\mathcal{J}A}$. Therefore, $O_B^{\mathcal{J}B} \subseteq \bar{\eta} \cup \{e\}$. Since $e \in R_B^{i_B}$, by Def. 2.30 it must be $e \notin O_B^{\mathcal{J}B}$ (because by Def. 2.8, $O^i \cap R^i = \emptyset$ for all i), hence $O_B^{\mathcal{J}B} \subseteq \bar{\eta}$. Since by hypothesis the payoff of B is zero in η , it must be $O_B^{\mathcal{J}B} \not\subseteq \bar{\eta}$ — contradiction. \square

B.5 Proof of Theorem 2.34

Assume that the statement (a) is true, i.e. all the participants agree on the contract $\mathcal{C} = \mathcal{C}_1 \mid \dots \mid \mathcal{C}_n$. By Def. 2.17, each $A_i \in \{A_1, \dots, A_k\} = \mathcal{A}$ has a winning strategy Σ_i in \mathcal{C} . Let σ be a fair play of \mathcal{C} conform to all the Σ_i . Since all the participants win in σ , by Lemma 2.33 there exists $k \in 1..n$ and a finite prefix η of σ such that $\Phi A_k \eta < 0$. By Def. 2.8, this amounts to say that there exists h such that $O_k^h \subseteq \bar{\eta}$ and $R_k^h \not\subseteq \bar{\eta}$.

We now prove that \mathcal{C}_k does *not* protect A_k . To do that, we construct a contract $\mathcal{C}' = \langle \mathcal{E}', \mathcal{A}, \pi', \Phi' \rangle$ such that A_k does not weakly agree on $\mathcal{C}' \mid \mathcal{C}_k$. By Def. 2.27, this amounts to proving that A_k has no non-losing strategies in $\mathcal{C}' \mid \mathcal{C}_k$. The functions π' and Φ' in \mathcal{C}' are almost immaterial: we just require that they make \mathcal{C}' compatible with \mathcal{C}_k . The ES \mathcal{E}' comprises a set of events \bar{E} disjoint from the events of \mathcal{C} and such that $\bar{e} \in \bar{E}$ for all event e in \mathcal{C} . The enablings and conflicts of \mathcal{E}' are the following:

$$\begin{aligned} & \{ \vdash e \mid e \in \bar{\eta} \setminus \pi^{-1}(A_k) \} \cup \{ \vdash \bar{e} \mid e \in \pi^{-1}(A_k) \} \\ & \{ e \# \bar{e} \mid e \in \pi^{-1}(A_k) \setminus \bar{\eta} \} \end{aligned}$$

Intuitively, \mathcal{C}' enables all the events of $\mathcal{A} \setminus \{A_k\}$ in η , and also all the events \bar{e} in conflict with some event of A_k , *except* for the events in η . The goal of \mathcal{C}' is to force A_k to do O_k^h , and then stop before R_k^h is reached. To implement this goal, \mathcal{C}' must also be innocent in η .

By contradiction, assume that Σ is a non-losing strategy for A_k in $\mathcal{C}' \mid \mathcal{C}_k$. Assume that \mathcal{C}' adopts the eager strategy $\Sigma_{\mathcal{C}'}$, i.e. all the enabled events are in its strategy. By fairness, there must exist a fair play ν of $\mathcal{C}' \mid \mathcal{C}_k$ which conforms to Σ and $\Sigma_{\mathcal{C}'}$, and where (1) all the participants are innocent (recall that the eager strategy is innocent, by Lemma 2.15) and (2) all the events \bar{e} in conflict with the events of A_k not in η have been performed before any other event of A_k not in η . By (1), A_k is innocent in ν , hence by Def. 2.14:

$$\forall i \geq 0. \forall e \in \pi^{-1}(A_k). (\bar{\nu}_i \xrightarrow{e} \implies \exists j \geq i. e_j \# e \vee e_j = e) \quad (8)$$

By (2), there cannot be events in conflict with those in $\bar{\eta} \cap \pi^{-1}(A_k)$. We can then rewrite equation (8) as follows:

$$\forall i \geq 0. \forall e \in O_k^h. (\bar{\nu}_i \xrightarrow{e} \implies \exists j \geq i. e_j = e)$$

Summing up, there exists a fair play ν of $\mathcal{C}' \mid \mathcal{C}_k$ which conform to Σ , and where $O_k^h \subseteq \bar{\nu}$. Furthermore, $R_k^h \not\subseteq \bar{\nu}$, because the ES \mathcal{E}' only enables the events in $\bar{\eta}$, while $\bar{\eta}$ does not contain all the events in R_k^h . By Def. 2.8, $\mathcal{W}A_k\nu = \Phi A_k\nu < 0$, i.e. A_k loses in ν — contradiction. \square

C Basic properties of CES

We define the predicate $DF(\sigma)$ true whenever σ has no duplicates, i.e. for all $i \leq n$, $e_i \notin \bar{\sigma}_i$.

Definition C.1 (Trace). Let $\mathcal{E} = (E, \#, \vdash, \Vdash)$ be a CES. A sequence $\sigma = \langle e_0 \dots e_n \rangle \in E^*$ is an X -trace of \mathcal{E} iff:

$$CF(\bar{\sigma}) \wedge DF(\sigma) \wedge \forall i \leq n. (e_i \in X \vee \bar{\sigma}_i \vdash e_i \vee \bar{\sigma} \Vdash e_i) \quad (9)$$

The set of all X -traces of \mathcal{E} is denoted by $\mathcal{T}_{\mathcal{E}}(X)$, abbreviated as $\mathcal{T}_{\mathcal{E}}$ when $X = \emptyset$.

For all sequences σ , we denote with $\sigma \downarrow$ the sequence obtained by eliminating from σ all the duplicate events. This is formalised in Def. C.2 below.

Definition C.2. For all $\sigma \in E^*$, we inductively define the sequence $\sigma \downarrow$ as follows:

$$\sigma \downarrow = \begin{cases} \varepsilon & \text{if } \sigma = \varepsilon \\ \sigma' \downarrow & \text{if } \sigma = \sigma' a \text{ and } a \in \bar{\sigma}' \\ \sigma' \downarrow a & \text{if } \sigma = \sigma' a \text{ and } a \notin \bar{\sigma}' \end{cases}$$

Lemma C.1. For all $X, Y \subseteq E$, and for all $\sigma, \eta \in E^*$:

- (a) $\sigma \in \mathcal{T}(\bar{\sigma})$
- (b) $X \subseteq Y \implies \mathcal{T}(X) \subseteq \mathcal{T}(Y)$
- (c) $\sigma \in \mathcal{T}(X) \wedge \sigma' \in \mathcal{T}(X) \wedge CF(\overline{\sigma\sigma'}) \implies (\sigma\sigma') \downarrow \in \mathcal{T}(X)$

Lemma C.2. For all $C, X \subseteq E$:

$$C \in \mathcal{F}(X) \iff \forall C_0 \subseteq_{fin} C. \exists \sigma \in \mathcal{T}(X). C_0 \subseteq \bar{\sigma} \subseteq C$$

Proof. (\implies) Let $C \in \mathcal{F}(X)$, and let $C_0 \subseteq_{fin} C$. By Def. 3.2, we see that $CF(C)$, and:

$$\forall e \in C_0. \exists \sigma^e \in \mathcal{T}(X). e \in \bar{\sigma}^e \subseteq C$$

Since C_0 is finite, we can concatenate all the (finite, conflict-free) sequences σ^e obtained above. Let σ' be the result of such operation. By C.1(c), $\sigma = \sigma' \downarrow \in \mathcal{T}(X)$. Also, by construction we have that $C_0 \subseteq \bar{\sigma} \subseteq C$, from which the thesis follows.

(\impliedby) Assume that $\forall C_0 \subseteq_{fin} C. \exists \sigma \in \mathcal{T}(X). C_0 \subseteq \bar{\sigma} \subseteq C$, and let $e \in C$. By the hypothesis, since $\{e\} \subseteq C$, there exists $\sigma^e \in \mathcal{T}(X)$ such that $e \in \bar{\sigma}^e \subseteq C$. It remains to prove that $CF(C)$. By contradiction, assume that $\neg CF(C)$. Then, there would exist $C_0 \subseteq_{fin} C$ such that $\neg CF(C_0)$, and so by hypothesis there would also exist some $\sigma \in \mathcal{T}(X)$ with $C_0 \subseteq \bar{\sigma}$. By Def. C.1, it must be $CF(\bar{\sigma})$, which contradicts $\neg CF(C_0)$. \square

The following result is a simple corollary of Lemma C.2. In case the set C is finite, to test if C is an X -configuration it suffices to find an X -trace which covers all the events of C .

Lemma C.3. *Let $C \subseteq_{fin} E$ be a finite set of events. We have that $C \in \mathcal{F}(X)$ iff there exists $\sigma \in \mathcal{T}(X)$ such that $\bar{\sigma} = C$.*

Lemma C.4. *For all $C, C', X, X' \subseteq E$, and for all $\sigma \in E^*$:*

$$(a) (C, X) \xrightarrow{\sigma} (C', X') \implies \bar{\sigma} = C' \setminus C \wedge C \cup \bar{\sigma} \Vdash X \setminus X'$$

$$(b) \sigma \in \mathcal{T}(X) \wedge X \text{ least credit for } \sigma \iff (\emptyset, \emptyset) \xrightarrow{\sigma} (C, X)$$

Definition C.3 (Trace with past). *For all $C \subseteq_{fin} E$, $X \subseteq E$, we say that $\sigma = \langle e_0 \dots e_n \rangle \in E^*$ is an X -trace with past C iff $CF(C \cup \bar{\sigma})$, $C \cap \bar{\sigma} = \emptyset$, $DF(\sigma)$, and*

$$\forall i \leq n. (e_i \in X \vee C \cup \bar{\sigma}_i \vdash e_i \vee C \cup \bar{\sigma} \Vdash e_i) \quad (10)$$

We denote with $\mathcal{T}_E^C(X)$ the set of X -traces with past C .

Lemma C.5. *Let (C, X) be a reachable state of LTS_E . Then:*

$$\sigma \in \mathcal{T}^C \iff \exists X_0 \subseteq X. (C, X) \xrightarrow{\sigma} (C \cup \bar{\sigma}, X_0)$$

Definition C.1. *For all $P, X \subseteq E$ we define the set $\mathcal{F}_E^P(X)$ of X -configurations with past P as follows:*

$$\mathcal{F}_E^P(X) = \{C \subseteq E \setminus P \mid CF(C) \wedge \forall e \in C. \exists \sigma \in \mathcal{T}_E^P(X). e \in \bar{\sigma} \subseteq C\}$$

For all $C, X \subseteq E$, the set of reachable events with past C and credits X is:

$$\mathcal{R}_E^C(X) = \bigcup \mathcal{F}_E^C(X)$$

C.1 Proof of Theorem 3.4

For coherence, let $\mathcal{A} \subseteq \mathcal{F}(X)$ be pairwise compatible in $\mathcal{F}(X)$. By definition of pairwise compatibility:

$$\forall e, e' \in \bigcup \mathcal{A}. \exists C \in \mathcal{F}(X). e, e' \in C$$

Since $C \in \mathcal{F}(X)$ implies $CF(C)$, it follows that $\neg(e \# e')$ for all $e, e' \in \mathcal{A}$, and so $CF(\bigcup \mathcal{A})$. Let $e \in \bigcup \mathcal{A}$. Then, there exists $C \in \mathcal{A}$ such that $e \in C$. Since $C \in \mathcal{F}(X)$, by Def. 3.2 there exists $\sigma \in \mathcal{T}(X)$ such that $e \in \bar{\sigma} \subseteq_{fin} C$. Since $C \subseteq \bigcup \mathcal{A}$, by Def. 3.2 we can conclude that $\bigcup \mathcal{A} \in \mathcal{F}(X)$.

Finiteness is straightforward by Def. 3.2, since for all $e \in C \in \mathcal{F}(X)$, the set of elements of the (finite) sequence $\sigma \in \mathcal{T}(X)$ such that $e \in \bar{\sigma} \subseteq_{fin} C$ is a configuration in $\mathcal{F}(X)$. \square

C.2 Proof of Theorem 3.6

(\Rightarrow) Let $C \in \mathcal{F}(X)$, and let $D \subseteq_{fin} C$. By Lemma C.2, there exists $\sigma \in \mathcal{J}(X)$ such that $D \subseteq \bar{\sigma} \subseteq C$. Let $X_0 \subseteq_{fin} X$ be the least credit for σ . By Lemma C.4, we have $(\emptyset, \emptyset) \xrightarrow{\sigma} (\bar{\sigma}, X_0)$. Therefore, the thesis follows by choosing $C_0 = \bar{\sigma}$.

(\Leftarrow) Let $D \subseteq_{fin} C$, and assume that $(\emptyset, \emptyset) \xrightarrow{\sigma} (C_0, X_0)$, for some σ and X_0 such that $C_0 = \bar{\sigma}$, $D \subseteq C_0 \subseteq C$, and $X_0 \subseteq X$. Assume that the trace has the form:

$$(\emptyset, \emptyset) \xrightarrow{e_0} (\bar{\sigma}_1, Y_1) \xrightarrow{e_1} (\bar{\sigma}_2, Y_2) \xrightarrow{e_2} \dots \xrightarrow{e_n} (\bar{\sigma}, Y_n)$$

where $Y_n = X_0$. By Lemma C.4, we have $\sigma \in \mathcal{J}(X_0)$ and since $X_0 \subseteq_{fin} X$, by Lemma C.1(b), it is also true that $\bar{\sigma} \in \mathcal{J}(X)$. Therefore, by Lemma C.2 we conclude that $C \in \mathcal{F}(X)$. \square

D Proofs for Section 3

D.1 Proof of Lemma 3.9

Straightforward from Def. 3.7, by taking $k = |\sigma| + 1$. Indeed, note that $e \notin X_k^{\sigma'}$.

D.2 Proof of Theorem 3.13

(\Rightarrow). We exploit the coinduction proof principle, by defining a functor F such that $\text{gfp } F$ is the coinductively defined set of prudent events. The coinduction proof principle states that:

$$x \sqsubseteq F(x) \implies x \sqsubseteq \text{gfp } F$$

Thus, we will obtain the thesis if, taking for x the set $\mathcal{R}_{\bar{\sigma}}^{\bar{\sigma}}$ of reachable events from $\bar{\sigma}$, we manage to prove that $x \sqsubseteq F(x)$.

More precisely, to fit the coinductive schema we consider an endofunctor F with the following type:

$$F : (E \times E^*) \times (\mathcal{A} \times E^*) \rightarrow (E \times E^*) \times (\mathcal{A} \times E^*)$$

First, define the predicate ϕ as follows:

$$\phi(\sigma, \sigma', \mathbf{A}) \triangleq \exists k > |\sigma|. X_k^{\sigma'} \cap \pi^{-1}(\mathbf{A}) \subseteq X_{|\sigma|}^{\sigma}$$

Then, for $P \subseteq E \times E^*$ and $I \subseteq \mathcal{A} \times E^*$, we define $F(P, I) = (P', I')$ as follows:

$$P' = \{(e, \sigma) \mid \exists \Sigma. \forall \sigma' = \sigma e \eta \text{ conform to } \Sigma. \forall \mathbf{B} \neq \pi(e). (\mathbf{B}, \sigma) \in I \implies \phi(\sigma', \sigma, \mathbf{A})\}$$

$$I' = \{(\mathbf{A}, \sigma) \mid \forall e \in \pi^{-1}(\mathbf{A}). \forall i \geq 0. \exists j \geq i. (e, \sigma_j) \notin P\}$$

To apply the coinduction proof principle, let:

$$\begin{aligned} P &= \{(e, \sigma) \mid e \in \mathcal{R}_{\bar{\sigma}}^{\bar{\sigma}}\} \\ I &= \{(\mathbf{A}, \sigma) \mid \forall e \in \mathcal{R}_{\bar{\sigma}}^{\bar{\sigma}} \cap \pi^{-1}(\mathbf{A}). e \in \bar{\sigma}\} \end{aligned}$$

The proof of $(P, I) \sqsubseteq F(P, I)$ proceeds then smoothly. The choice of the strategy Σ in P' is made as follows. For all η and for all $\mathbf{B} \in \mathcal{E}$, let:

$$\Sigma_{\mathbf{A}}(\eta) = \mathcal{R}^{\bar{\eta}} \cap \pi^{-1}(\mathbf{A})$$

Notice that, since \mathcal{E} is conflict-free, then all these strategies are well-defined.

(\Leftarrow). Assume that e is prudent for \mathbf{A} in σ . We must prove that $e \in \mathcal{R}^{\bar{\sigma}}$. For all participants \mathbf{B} , consider the greatest prudent strategy $\Sigma_{\mathbf{B}}^p$. Let $\sigma' = \sigma e \nu$ be such that ν conforms to *all* the strategies $\Sigma_{\mathbf{B}}$. By Lemma 3.12, all participants are innocent in σ' . Since all events in $e \nu$ are prudent, then by Def. 3.7:

$$\forall \mathbf{B}. \exists k_{\mathbf{B}} > |\sigma|. X_{k_{\mathbf{B}}}^{\sigma'} \cap \pi^{-1}(\mathbf{B}) \subseteq X_{|\sigma|}^{\sigma'}$$

By choosing $k = \max \{k_{\mathbf{B}} \mid \mathbf{B} \in \mathcal{A}\}$:

$$\exists k > |\sigma|. X_k^{\sigma'} \subseteq X_{|\sigma|}^{\sigma'}$$

By Lemma C.5, we have $e \nu \in \mathcal{T}^{\bar{\sigma}}$. Thus, by Def. C.1, we conclude that $e \in \mathcal{R}^{\bar{\sigma}}$. \square

D.3 Proof of Lemma 3.19

By Def. 3.18, the CES $\mathcal{E}(\Phi)$ is conflict-free, and its minimal enablings have the form $R_{\mathbf{A}}^i \Vdash O_{\mathbf{A}}^i$, for all $\mathbf{A} \in \mathcal{A}$ and for all i . Consider the set C containing all and only the events in the minimal enablings of $\mathcal{E}(\Phi)$. Since Φ has a finite number of offers-requests, and each of them is finite, then C is finite. We shall prove that $C \in \mathcal{F}_{\mathcal{E}(\Phi)}$. Since C is finite, by Lemma C.3 it suffices to find some trace $\sigma \in \mathcal{T}_{\mathcal{E}(\Phi)}$ such that $\bar{\sigma} = C$. Consider an arbitrary ordering σ of the events of C . For all $e \in \bar{\sigma}$, we now prove that $\bar{\sigma} \Vdash e$. By Def. 3.2, this will allow to conclude that σ is a trace of $\mathcal{E}(\Phi)$. Let $e \in \bar{\sigma}$. We have the following two cases:

- $e \in O_{\mathbf{A}}^i$, for some i and $\mathbf{A} \in \mathcal{A}$. Since $R_{\mathbf{A}}^i \Vdash O_{\mathbf{A}}^i \in \mathcal{E}(\Phi)$ and all events in $R_{\mathbf{A}}^i$ occur in σ , then $\bar{\sigma} \Vdash e$.
- $e \in R_{\mathbf{A}}^i$, for some i and $\mathbf{A} \in \mathcal{A}$. By circularity (Def. 2.30), there exist j and \mathbf{B} such that $e \in O_{\mathbf{B}}^j$. By Def. 3.18, $R_{\mathbf{B}}^j \Vdash O_{\mathbf{B}}^j \in \mathcal{E}(\Phi)$, and the proof proceeds as in the previous case.

Summing up, we have found that $C \in \mathcal{F}_{\mathcal{E}(\Phi)}$. By construction, C contains all the requests $R_{\mathbf{A}}^i$, for all i and \mathbf{A} . \square

D.4 Proof of Theorem 3.20

Let \mathcal{C} be a contract with O-R payoffs for \mathbf{A} . Let $C \in \mathcal{F}_{\mathcal{E}}$ be such that $\bigcup_i R_{\mathbf{A}}^i \subseteq C$. We will prove that the prudent strategy $\Sigma_{\mathbf{A}}^p$ is winning for \mathbf{A} in \mathcal{C} . Let γ be a fair play of \mathcal{C} which conforms to $\Sigma_{\mathbf{A}}^p$.

By contradiction, assume that \mathbf{A} is *not* winning in γ . By Lemma 3.12, \mathbf{A} is innocent in γ . By Lemma 3.15, \mathbf{A} is credit-free in γ . Thus, by Def. 3.14 it

follows that all participants are innocent and $\Phi A \gamma \leq 0$. By Def. 2.8, this means that either there exists some i such that $O_A^i \subseteq \bar{\gamma}$ and $R_A^i \not\subseteq \bar{\gamma}$ (in case A loses), or for all i , $O_A^i \not\subseteq \bar{\gamma}$ and $R_A^i \not\subseteq \bar{\gamma}$ (in case A non-loses). In both cases, there exists at least one i such that $R_A^i \not\subseteq \bar{\gamma}$.

Let i be such that $R_A^i \not\subseteq \bar{\gamma}$, and let e be such that $e \in R_A^i \setminus \bar{\gamma}$. By hypothesis, there exists $C \in \mathcal{F}_\mathcal{E}$ such that $\bigcup_i R_A^i \subseteq C$, hence $e \in C$. Since C is a configuration, it enjoys finiteness (Theorem 3.4), i.e. there exists $C_0 \subseteq_{fin} C$ such that $e \in C_0 \in \mathcal{F}_\mathcal{E}$. Since \mathcal{E} is \vdash -free, it must be $C_0 \Vdash C_0$, i.e. all the events in C_0 are circularly enabled by C_0 itself.

We will prove that $C_0 \subseteq \bar{\gamma}$. Let $\gamma = \langle e_0 e_1 \dots \rangle$. Since all participants are innocent in γ , by Theorem 3.13 it must be:

$$\forall i \geq 0. \forall e. (e \in \mathcal{R}_\mathcal{E}^{\sigma_i} \implies \exists j \geq i. e_j = e) \quad (11)$$

Now, since $e \in C_0 \in \mathcal{F}_\mathcal{E}$, by definition of reachable events we have $e \in \mathcal{R}_\mathcal{E}^\emptyset$, hence by (11) it follows that $e \in \bar{\gamma}$ — contradiction.

Summing up, we have proved that $\bigcup_i R_A^i \subseteq \bar{\gamma}$ for all fair plays γ . Therefore, A has a winning strategy (Σ_A^p) in \mathcal{C} , and so we conclude that A agrees on \mathcal{C} . \square

D.5 Proof of Theorem 3.21

Let \mathcal{C}_A be a contract with O-R payoffs for A such that:

$$\forall i, Y. Y \triangleright O_A^i \implies R_A^i \subseteq Y \quad (12)$$

where we write $Y \triangleright O_A^i$ as a shorthand for:

$$\forall e \in O_A^i. Y \vdash e \vee Y \Vdash e$$

Let \mathcal{E} be the CES of contract \mathcal{C}_A , and let \mathcal{C} be a contract compatible with \mathcal{C}_A . Choose for A the strategy Σ_A^p which enables all and only the prudent events. To prove that \mathcal{C}_A protects A, we will prove that Σ_A^p is a non-losing strategy for A in $\mathcal{C}_A \mid \mathcal{C}$. Let σ be a fair play of $\mathcal{C}_A \mid \mathcal{C}$ conform to Σ_A^p . By Lemma 3.12, A is innocent in σ . By contradiction, assume that A loses in σ , i.e. by Def. 2.16 and by Def. 2.8:

$$\exists i. O_A^i \subseteq \bar{\sigma} \wedge R_A^i \not\subseteq \bar{\sigma} \quad (13)$$

Since \mathcal{E} is finite, then by Lemma 3.11 the strategy Σ_A^p is prudent. Let $\sigma = \langle e_0 e_1 \dots \rangle$, and let $e_j \in O_A^i$. Since e_j is prudent in σ_j , it must have been enabled by a \vdash or by a \Vdash . There are the following two cases:

- $Z_j \vdash e_j$, for some $Z_j \subseteq \bar{\sigma}_j$. Then, $\bar{\sigma} \vdash e_j$.
- $Z_j \Vdash e_j$, for some Z_j . Since Σ_A^p is prudent, then by Lemma 3.15 A is credit-free in σ , and so $Z_j \subseteq \bar{\sigma}$.

Summing up, $\bigcup\{Z_j \mid e_j \in O_A^i\} \subseteq \bar{\sigma} \triangleright O_A^i$. Therefore, by (12) it follows that $R_A^i \subseteq \bar{\sigma}$ — which contradicts (13). \square

D.6 Proof of Theorem 3.22

Let Φ_1, \dots, Φ_n be finite circular O-R payoffs for participants A_1, \dots, A_n , respectively, such that, for all $A \in \mathcal{A} = \{A_1, \dots, A_n\}$ condition (3.22) holds, i.e.:

$$\forall P \subseteq \mathbb{N}. \forall j. O_A^j \subseteq \bigcup_{i \in P} O_A^i \implies R_A^j \subseteq \bigcup_{i \in P} R_A^i$$

For all $i \in 1..n$, let $\mathcal{E}_i = \mathcal{E}(\Phi_i)$ be the CES of contract \mathcal{C}_i . Let $\mathcal{C} = \mathcal{C}_1 \mid \dots \mid \mathcal{C}_n$, and let $\Phi = \Phi_1 \sqcup \dots \sqcup \Phi_n$. By construction, the CES of \mathcal{C} is $\mathcal{E}(\Phi) = \mathcal{E}_1 \sqcup \dots \sqcup \mathcal{E}_n$. We prove the three items of Theorem 3.22 separately:

(a) By Lemma 3.19, we have that:

$$\exists C \in \mathcal{F}_{\mathcal{E}(\Phi)}. \forall A_p \in \mathcal{A}. \bigcup_i R_{A_p}^i \subseteq C$$

Since $\mathcal{E}(\Phi)$ is conflict-free and \vdash -free, by Theorem 3.20 it follows that each A_i agrees with \mathcal{C} . Therefore, \mathcal{C} admits an agreement.

(b) We prove that \mathcal{C}_i protects A_i by exploiting Theorem 3.21. Denote the offers-requests of A_i just as $(O^j)_j$ and $(R^j)_j$, respectively. Let Y be such that

$$Y \subseteq \bigcup \{X \mid X \Vdash e \in \mathcal{E}_i \wedge e \in O^j\} \quad \text{and} \quad Y \Vdash O^j$$

Let $O^j = \{e_1, \dots, e_p\}$. Recall that $Y \Vdash O^j$ stands for $\forall h \in 1..p. Y \Vdash e_h$. By construction of $\mathcal{E}(\Phi_i)$:

$$\forall h \in \{1..p\}. \exists i_h. e_h \in O^{i_h} \wedge R^{i_h} \subseteq Y$$

Thus, $O^j \subseteq \bigcup_{k \in 1..p} O^{i_k}$. Since condition (5) holds by hypothesis, it must be:

$$R^j \subseteq \bigcup_{k \in 1..p} R^{i_k} \subseteq Y$$

We can then apply Theorem 3.21, and deduce that \mathcal{C}_i protects A_i .

(c) Let σ be a play of \mathcal{C} . Since plays are traces of $\text{LTS}_{\mathcal{E}(\Phi)}$, they can only contain events e such that $X \circ e \in \mathcal{E}(\Phi)$, for some X and $\circ \in \{\vdash, \Vdash\}$. By Def. 3.18, all enablings of $\mathcal{E}(\Phi)$ have the form $R_A^i \Vdash O_A^i$. Therefore, each event e in σ belongs to $O_{\pi(e)}^i$, for some i . \square