

Reversible and Composable Financial Contracts (extended abstract)

Vincent Danos, Jean Krivine and Julien Prat

Reversible and Composable Financial Contracts (extended abstract)

Vincent Danos, Jean Krivine, Julien Prat

► **To cite this version:**

Vincent Danos, Jean Krivine, Julien Prat. Reversible and Composable Financial Contracts (extended abstract). 2020. hal-03103298

HAL Id: hal-03103298

<https://hal.archives-ouvertes.fr/hal-03103298>

Preprint submitted on 8 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reversible and Composable Financial Contracts (extended abstract)

Vincent Danos (CNRS, École Normale Supérieure) Jean Krivine (CNRS, Université de Paris)

Julien Prat (CNRS, École Polytechnique)

October 25, 2020

Abstract

It is widely believed that financial markets cannot be liquid without centralised processes to manage counterparty risk. We propose an alternative method for liquidity based on reversible and composable contracts run atop a blockchain. Novel instruments for zero-collateral intermediation can be defined.

1 Introduction

Financial trades face a liquidity/risk trade-off [1]. Forcing real-time settlements introduces cash constraints for Buyers [3] since they must have the cash on hand before entering a trade. This precludes short-selling, impedes intermediation, ultimately hurting liquidity. One can relax the cash constraint by allowing for deferred payments. But deferred payments introduce risk. How does one make sure that a deferred payment is eventually made?

We introduce a trade protocol where payment defaults are handled by reversing trades. Not only can players buy without paying first, but, crucially, they can sell again while still withholding payments. Trades are no longer disjoint bilateral contracts resolved independently, but aggregate into *trade lines*. The ownership of an underlying asset becomes distributed among players with positions in the trade line. The course of the game determines who ends up owning that asset and the overall payoffs of the participants. The game can be implemented as a smart contract on a blockchain.

We also define and investigate a specific liquidity instrument (§7). Its basic brick is a standard cancellation contract: if Buyer has not paid the agreed amount by a certain date, Seller can cancel the deal and receive a compensation fee. By chaining two copies of this contract, one obtains a ternary trade line. The question is whether one can set the parameters of each bilateral contract (terms and fees) so that the middle player has an incentive to play the composite game. Indeed, we show that a middle player achieves zero-collateral intermediation in the game-theoretic equilibrium provided certain natural conditions on the contract parameters hold (and players are rational).

We start with the definition of unilateral contracts (§2) and associated operations. We then define bilateral contracts and give standard examples of such (§3). Then we define tradelines (§4) and the trade game proper. We establish quantitative soundness (§5) and stability properties of the trade game (§6). In (§7) we give a heuristic derivation of the game-theoretic equilibrium of a parameterised ternary instance of the trade game, and infer conditions on the parameters for the game to provide liquidity. Proofs are omitted to save space.

2 Unilateral contracts

Our first task is to set up a language for elementary contracts -which we call *clauses*. Throughout we use a functional notation for clauses and see them as functions of pairs of positions. A clause embodies a (time-dependent) promise of one player to another. It is directed. Bilateral contracts (next Section) consist of pairs of reciprocal promises.

We write $T \in \mathbb{Z}$ for the global time variable (block time in the basic implementation). We assume a countable set of positions.

Definition 1 A clause over positions x, y is defined as:

$$\begin{aligned}
\Theta xy &= \bigoplus_{I \in \mathcal{I}} T \in I \mapsto \theta xy && \text{time-dependent clause} \\
\theta xy &= \theta xy \vee \theta xy \mid \gamma xy \vdash \eta xy && \text{clause} \\
\gamma xy &= \top \mid \perp \mid \omega xy \mid \gamma xy \wedge \gamma xy && \text{guard} \\
\omega xy &= (x \rightarrow^a y)_+ \mid (y \rightarrow^a x)_- && \text{active payment} \\
\eta xy &= y \rightarrow^b x \mid \eta xy; \eta xy \mid \varepsilon && \text{passive payment (effect)}
\end{aligned}$$

with \mathcal{I} a finite family of disjoint time intervals, $a \in \mathbb{R}^+$, $b \in \mathbb{R}$.

Thus a clause is a time-dependent disjunction of pairs of guards and effects.

In a guard γxy , players at positions x and y play dual roles: x 's player can activate the clause with a payment $(x \rightarrow^a y)_+$, while y 's can inhibit the clause with a payment $(x \leftarrow^a y)_-$. One says that x is the *active* position in γxy , while y is the *reactive* one.

Payments included in guards are referred to as *active payments*. Clauses also incorporate *passive payments* or effects $\eta xy = y \rightarrow^b x$, with $b \geq 0$ when payment flows from y to x , ie from the reactive to the active position. The direction of passive payments is unconstrained. As their name indicate, passive payments are made as a consequence of a (forward or backward) move being played and require no intervention from the players.

We write $\mathbf{1} := \top \vdash \varepsilon$ for the trivial clause with a guard which is always true and has no effect.

We define the following quotient over payment expressions of the same type:

$$\begin{aligned}
(x \rightarrow^0 y)_\varepsilon &= \top \\
(x \rightarrow^a y)_\varepsilon \wedge (x \rightarrow^{a'} y)_\varepsilon &= (x \rightarrow^{a+a'} y)_\varepsilon
\end{aligned}$$

whereby any guard γxy which is not \perp can be rewritten uniquely to a simple conjunction $(x \rightarrow^a y)_+ \wedge (x \rightarrow^{a'} y)_-$ with $a, a' \geq 0$.

One can identify a guard with a pair a, a' of non-negative real numbers, and a real number b describing its effect part. Thus, one has a concrete form for clauses as finite sets of tuples (a, a', b) with $a \geq 0$ an (active) activation payment, $a' \geq 0$ an (active) inhibition payment, and b a passive effect payment. Each element of the set corresponds to one disjunct of the clause disjunction. In this *linear* representation a clause is simply an element $\mathcal{P}_{\text{fin}}(\mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R})$. For example $\mathbf{1}$ is $(0, 0, 0)$.

► Hereafter, we will mostly assume *deterministic* clauses, where the disjunction has just one term. This keeps notations lighter, while the generalisation to non-deterministic clauses is straightforward. To simplify things further, we leave out inhibition payments. However we will use and discuss both extensions freely in examples.

To express interesting contracts, one needs time-dependent clauses. They are defined as partial and finite piecewise-constant functions of discrete time with values in (constant) clauses. We will write $\Theta(t)$ when we want to stress time-dependency.

The latest finite date mentioned in Θ is called its *time horizon*, and the set of times at which Θ is defined $|\Theta| = \cup \mathcal{I}$, is called its *time domain*. Outside its time domain, Θ is undefined and cannot be

triggered. Past its horizon, its guards and effects can no longer change if they are defined. One says Θ is *eventually defined* if $[M, +\infty) \subseteq |\Theta|$ for some M .

Using the linear representation of clauses, time-dependent clauses can be seen as finitary piece-wise constant and partial maps from \mathbb{Z} to $\mathcal{P}_{\text{fin}}(\mathbb{R}^2)$ (or \mathbb{R}^3 if one includes inhibition payments). We write $\mathcal{C} \subseteq \mathbb{Z} \rightarrow \mathcal{P}_{\text{fin}}(\mathbb{R}^2)$ for the set of such clauses. Wherever a time-dependent clause Θ is undefined, its linear representation returns the empty set. Therefore, one can redefine equivalently in the linear representation the domain of a clause as $|\Theta| = \{t \mid \Theta(t) \neq \emptyset\}$, and the set of deterministic clauses \mathcal{C}_0 as those such that $\Theta(t)$ has always at most one element.

We will use onwards whichever representation is most convenient for the question at hand.

Definition 2 Let Θ, Θ' be in \mathcal{C} , define:

$$(\Theta + \Theta')(t) = \Theta(t) + \Theta'(t) \quad (1)$$

$$(\Theta \vee \Theta')(t) = \Theta(t) \vee \Theta'(t) \quad (2)$$

In Eq. 1, the rhs is the pairwise sum $\{\theta_i\} + \{\theta'_j\} = \{\theta_i + \theta'_j\}$; in Eq. 2, the rhs is the union. If either $\Theta(t)$ or $\Theta'(t)$ is the empty set, then so is the product $\Theta(t) \times \Theta'(t)$ and therefore $(\Theta + \Theta')(t) = \emptyset$ as well. Hence $|\Theta + \Theta'| = |\Theta| \cap |\Theta'|$. Clearly, $|\Theta \vee \Theta'| = |\Theta| \cup |\Theta'|$.

We will use the additive operator to shift clauses from one contract to another when composing contracts. Shifting clauses may lead to clauses with smaller domains.

Clause addition $\Theta + \Theta'$ should not be confused with the disjoint sum $\Theta \oplus \Theta'$ (used in def. 1) which is a particular case of union.

Proposition 1 The triple $(\mathcal{C}, \vee, +)$ form a commutative idempotent semi-ring, with respective neutral elements: the everywhere defined clause with value $\mathbf{1}$, and the nowhere defined clause $\mathbf{0}$. The subset \mathcal{C}_e of eventually defined clauses form a sub-semi-ring of \mathcal{C} . The domain map is a semi-ring morphism from $(\mathcal{C}, \vee, +, \mathbf{0}, \mathbf{1})$ to $(\mathcal{P}_{\text{fin}}(\mathbb{Z}), \cup, \cap, \emptyset, \mathbb{Z})$.

The union operator may lead in general to non-deterministic clauses, but absent inhibition payments, it has a compelling determinisation. The idea is that, at any given time, the active player can compare her options and pick up the most rewarding -ie minimise expenses and amount provisioned. This partial ordering can be linearised by maximising funds locked in the protocol, leading to:

Definition 3 Let Θ, Θ' be in \mathcal{C} , define:

$$\max(\Theta, \Theta')(t) = \max(\Theta(t) \vee \Theta'(t)) \quad (3)$$

As for the plain union we have $|\max(\Theta, \Theta')| = |\Theta| \cup |\Theta'|$. The novelty is that the determinised union \max preserves deterministic clauses, just like the addition does.

The subset of deterministic clauses \mathcal{C}_0 also forms a semi-ring $(\mathcal{C}_0, \max, +)$; the unary version of \max , \max_1 is a semi-ring morphism from $(\mathcal{C}, \vee, +)$ to $(\mathcal{C}_0, \max, +)$; it is onto and preserves domains in the sense that $|\cdot| \circ \max_1 = |\cdot|$.

3 Bilateral contracts

Definition 4 A bilateral contract, or trade arc, consists of a pair of time-dependent clauses ϕ, β , and a pair of positions: u, v ; where ϕ, β are called the forward and backward clauses, u, v are called the Seller's and Buyer's position. We use the following graphical notation to summarise the data:

$$u \begin{array}{c} \xrightarrow{\phi v u} \\ \xleftarrow{\beta u v} \end{array} v$$




Figure 1: Tree of the basic game: ϵ is a parameter representing the constant cost of a move; π_B (π_F) is the aggregated payoff to u (v) described in clause βuv (ϕvu); the backward (forward) move B_u (F_v) is only available in contexts where βuv (ϕvu) holds.

The forward clause specifies conditions under which Buyer can complete the trade, and conditions under which Seller can block its completion. Symmetrically, the backward clause specifies conditions under which Seller can cancel the trade, and conditions under which Seller can block its cancellation. The basic game is symmetric. The Seller/Buyer distinction only makes sense when basic games are composed, and there is a distinguished root position which holds an asset underlying the trade line (§4).

Note that, in the forward clause, the active position is Buyer's, that is to say the clause is ϕvu (not ϕuv). In the backward one, the active position is Seller's, and the clause is βuv .

Clauses are evaluated in a context which includes the current time and a complete list of (active) payments made. We do not describe explicitly the part of the context tracking payments. A trade arc ϕ , β between u and v is said to be in an F -state (B -state) in a given context, if ϕvu (βuv) holds in this context. Buyer (Seller) can play his move and complete (cancel) the trade iff the arc is in an F -state (B -state).

Notice that the game tree (Fig. 1) is not strictly speaking that of a sequential game [2]: the availability of moves to either player is context-dependent and so are the payoffs; and moves may (depending on the context) be available simultaneously to both players.

Definition 5 A bilateral contract ϕ , β is said to be *BF-exclusive* (or simply *exclusive*) if $|\phi| \cap |\beta| = \emptyset$. It is said to be *idle* in a given context, if it is neither in a B -state or an F -state in that context. A bilateral contract ϕ , θ is said to be *eventually-F* if ϕ is eventually defined, and *eventually-B* if β is.

If a contract is *BF-exclusive* there is no context in which it can be *both* in a B -state and an F -state. Absent this property, players may move simultaneously. In a blockchain implementation, this means that plays depend on the order in which the moves are ordered by the block-makers -clearly not a good thing.

To avoid deadlocks, one can restrict contracts to be eventually never idle. This way no player can lock a contract inadvertently.

Definition 6 The standard bilateral contract is defined as:

$$\begin{aligned} \phi vu &= T \geq -\infty \mapsto (v \rightarrow^a u)_+ \vdash \epsilon \\ \beta uv &= T \geq \Delta \mapsto \top \vdash v \rightarrow^p u \end{aligned}$$

where Δ is called the delay, a the price, and $p \geq 0$ the penalty.

Buyer has an option to conclude the deal at any time, and thus obtain some underlying asset irreversibly by paying a to Seller. However, starting at $T = \Delta$, and onwards, Seller has the right (but not the obligation) to cancel the deal and recall the underlying asset. When that happens Buyer has to pay a penalty p to Seller for immobilising the asset. This contract defined above is *eventually-B*, as after its time horizon Δ , the β -clause holds forever.

One can refine it by introducing another delay Δ' to restrict the forward clause:

$$\phi'vu = T < \Delta' \mapsto (v \rightarrow^a u)_+ \vdash \varepsilon$$

so that $|\phi'| = (-\infty, \Delta')$, $|\beta| = [\Delta, +\infty)$. Now Buyer has a limited time to exercise his buying option. The refined contract is BF -exclusive iff $\Delta' \leq \Delta$, and if $\Delta' < \Delta$, it will be idle during the $[\Delta', \Delta)$ interval. For $\Delta' = +\infty$, we recover the original contract.

4 Composite contracts

Definition 7 A tradeline is a non-empty list of positions u_1, \dots, u_n connected by trade arcs:

$$u_1 \xrightarrow[\beta_1]{\phi_1} u_2 \xrightarrow[\beta_2]{\phi_2} \dots u_{n-1} \xrightarrow[\beta_{n-1}]{\phi_{n-1}} u_n$$

u_1 is called the origin and u_n the end of the trade line. If $u_1 = u_n$, one says the trade line is resolved. If, in a given context: no trade arc is in an F - or a B -state, one says the trade line is irreducible, or a normal form (NF); if every arc is in an F - or a B -state, one says the trade line is connected.

Each position in a trade line has an owner. We often designate the owner of a position by the position itself, and speak of player u instead of the player owning u . A trade line represents the state of a game being played between the owners of its positions. Each move induces modifications of the trade line and payoffs to the players. Moves are attached to positions (and not directly to players). We distinguish two types of moves: *contraction moves* which generalise the moves already considered in the bilateral contracts, and *extension moves* whereby the trade line grows and new players come in the game.

4.1 Contractions

Contractions are defined as follows. The active player and clause are in red, and we assume that the active clause holds in the current context. The evicted player is in blue. Fusion operators μ_{XY} are defined below. The rest of the trade line before u , and after w stays unchanged. Effects attached to the trigger clause are evaluated at the same time as the transition is performed.

$$u \xrightarrow[\beta]{\phi} v \xrightarrow[\beta']{\phi'} w \xrightarrow[\phi vu]{F_v} v \xrightarrow[\mu_{FB}(\beta, \beta')]{\mu_{FF}(\phi, \phi')} w \quad (4)$$

$$u \xrightarrow[\beta]{\phi} v \xrightarrow[\beta']{\phi'} w \xrightarrow[\beta uv]{B_u} u \xrightarrow[\mu_{BB}(\beta, \beta')]{\mu_{BF}(\phi, \phi')} w \quad (5)$$

It may be that that v is the end of the trade line, in which case we write $v \cdot$:

$$u \xrightarrow[\beta]{\phi} v \cdot \xrightarrow[\phi vu]{F_v} v \cdot \quad (6)$$

$$u \xrightarrow[\beta]{\phi} v \cdot \xrightarrow[\beta uv]{B_u} u \cdot \quad (7)$$

If there is no position before u either, the trade line has fully resolved and the owner of the last remaining position is now in full possession of the underlying asset.

A connected trade line is one where each arc defines at least one contraction move; an irreducible trade line is one where no arc defines a contraction move.

Transitions (4) and (5) involve binary operations on clauses: μ_{FB} , μ_{FF} (forward fusions), and μ_{BB} , μ_{BF} (backward fusions). In game terms, one can think of a contraction as a move whereby the owner

of an active position acquires an adjacent position. Fusions give a precise meaning to ‘acquiring a new position’. Fusions are the core of *novation* (a new party substitutes for another in a contract) and decide the nature of the composite game being played.

Definition 8 (standard fusions) *The standard fusions are as follows:*

$$\begin{aligned}\mu_{FF}(\phi, \phi') &= \mu_{BF}(\phi, \phi') = \phi + \phi' \\ \mu_{FB}(\beta, \beta') &= \beta' \\ \mu_{BB}(\beta, \beta') &= \beta\end{aligned}$$

The idea behind the choice of μ_{BB} is that u is carrying over his original cancellation condition in the new contract with w (more about this shortly). The idea behind μ_{FB} is that the contraction preserves the cancellation condition of v . Finally, the idea behind the choice for the forward fusions μ_{FF} , and μ_{BF} , is that the payment just made (or promised) by v is transferred onto w 's forward clause (with clause sum as defined in Eq. 1).

Instead of the standard backward fusion $\mu_{BB}(\beta, \beta') = \beta$, we could keep only the logical part of β and not propagate its effects -eg penalties or fees imputed to the middle position v for not playing F_v in the standard contract (Def. 6). Another option is $\mu_{BB}(\beta, \beta') = \beta \vee \beta'$ (with clause union as defined in Eq. 2). All are reasonable in that the fused backward clause $\mu_{BB}(\beta, \beta')$ which ties in u and w is implied by β . This means that the player triggering the B -move does not have to stop after the first contraction and *can sweep the entire trade line* if she wishes to, for as long as β holds.

With the combinators obtained earlier, we can define an alternative to Def. 8:

Definition 9 (semi-ring fusions) *The semi-ring fusions are as follows:*

$$\begin{aligned}\mu_{FF}(\phi, \phi') &= \mu_{BF}(\phi, \phi') = \phi + \phi' \\ \mu_{BB}(\beta, \beta') &= \mu_{FB}(\beta, \beta') = \beta \vee \beta'\end{aligned}$$

With this choice, all calculations happen in the semi-ring structure. This set of fusions also satisfies the B -sweep property mentioned above.

Lemma 1 *Let γ, γ' be trade lines, and suppose γ contracts to γ' under semi-ring fusions. The following holds: (i) if γ is connected, so is γ' ; (ii) if γ is exclusive, so is γ' ; (iii) if all arcs in γ are eventually- B or $-F$, so are the ones in γ' .*

Standard fusions on the other hand satisfy only point (ii), and a weaker form of (iii). We will work with both sets of fusions.

4.2 Extension

A player at the end u of the trade line can extend it using the sell rule $S_{uv}(\phi, \beta)$ and append a new bilateral contract with clauses ϕ, β , thereby adding a new position v to the game:

$$u \cdot \xrightarrow{S_{uv}(\phi, \beta)} u \xrightarrow[\beta]{\phi} v \cdot \quad (8)$$

Contractions are under the control of a single player: Seller for the backward contraction, and Buyer for the forward one. Extensions are different. Both players must consent.

One may wonder whether extensions could be made at the start of the trade line as well (or even inserted in the middle of it). In fact, such extensions are dangerous. Suppose one extends a trade line by introducing a new position u_0 left of the origin:

$$u_0 \xrightarrow[\text{⊢-}u_1 \rightarrow^a u_0]{0} u_1 \longrightarrow \cdots \longrightarrow u_n$$

The new player at the origin u_0 can now sweep through the entire line by iterating B_{u_0} and collect an arbitrary fee a from everyone. In §5 we show that our design choices prevent such catastrophic events.

We can return to the standard contracts defined earlier (§6). With simplified and self-evident notations we get the following contraction rules using (under standard fusions):

$$\begin{aligned}
u \xrightarrow[\Delta_1, p_1]{\Delta'_1, a_1} v \xrightarrow[\Delta_2, p_2]{\Delta'_2, a_2} w &\xrightarrow[v \rightarrow^{p_1} u]{B_u} u \xrightarrow[\Delta_1, p_1]{\min(\Delta'_1, \Delta'_2), a_1 + a_2} w \quad \text{for } T \geq \Delta_1 \\
u \xrightarrow[\Delta_1, p_1]{\Delta'_1, a_1} v \xrightarrow[\Delta_2, p_2]{\Delta'_2, a_2} w &\xrightarrow[v \rightarrow^{a_1} u]{F_v} v \xrightarrow[\Delta_2, p_2]{\min(\Delta'_1, \Delta'_2), a_1 + a_2} w \quad \text{for } T < \Delta'_1
\end{aligned}$$

Note that the backward clauses depend on the previous move. To restore the symmetry, we could use the semi-ring fusion to obtain:

$$\beta x w = (T \geq \Delta_1 \mapsto \top \vdash w \rightarrow^{p_1} x) \vee (T \geq \Delta_2 \mapsto \top \vdash w \rightarrow^{p_2} x)$$

with $x = u, v$ depending on the preceding move. This means that at $\max(\Delta_1, \Delta_2)$ and onwards, the player in Seller position x can choose to be paid either penalties p_1 or p_2 , and will likely choose the highest. This is an instance of the determinising idea seen in Def. 3.

5 Soundness of the trade game

A basic form of soundness of the trade game is that it should never be possible to end up in a state where the asset underlying the trade is lost because the state cannot be resolved to a single position. As said, a way to solve the problem is to restrict to arcs which are eventually- B (or F). In this section, we are looking for a more quantitative notion of soundness, namely an upper bound on the expenses incurred by playing the game.

Such an upper bound also allows one to statically compute the amount of cash to stake in, upon entering the game, so that passive payments (in effects) are fully provisioned.

► The derivation assumes standard fusions and deterministic clauses. Similar bounds can be derived with the semi-ring ones and/or general clauses. To follow the calculations below it is convenient to take a player-centric view. Consider a generic position v in a trade line:

$$\dots \longrightarrow u \xrightarrow[\beta]{\phi} v \xrightarrow[\beta']{\phi'} w \longrightarrow \dots$$

Let t be the current time; we decompose the current values of the various clauses of interest in their components map (assuming they are all defined at t):

$$\begin{aligned}
\phi(t) &= \phi^1(t), \phi^2(t) & \beta(t) &= \beta^1(t), \beta^2(t) \\
\phi'(t) &= \phi'^1(t), \phi'^2(t) & \beta'(t) &= \beta'^1(t), \beta'^2(t)
\end{aligned}$$

These define four contraction moves which are the only moves in the trade line that impinge directly on v 's payoffs. Moves and payoffs are shown in Fig. 2. For F_v, B_v (solid lines) v is the active player; for B_u, F_w (dotted lines) v is passive and evicted by the move. Accordingly, from v 's viewpoint we call, $\phi^2(t) - \phi^1(t)$, and $\beta'^2(t) - \beta'^1(t)$ active payoffs, and $\beta^1(t) - \beta^2(t)$, $\phi'^1(t) - \phi'^2(t)$ passive ones.

Let γ be a trade line, and let t be a time. We denote by $<$ the positional ordering in γ . For any position $i \in \gamma$ not at the end of γ , we denote by $\phi_i(t), \beta_i(t)$ the clauses where i is Seller.

Payoffs at fixed time are specified by pairs of real numbers. We use the following notations: $\phi_i^1(t) \geq 0$ is the active payment i 's Buyer makes to i to complete the deal; and $\phi_i^2(t)$ is the (possibly negative) passive payment from i to his Buyer which follows. When i 's Buyer plays that forward move, i 's payoff is therefore: $\phi_i^1(t) - \phi_i^2(t)$.




Figure 2: Moves which change the balance of v and the implied payoffs.

► We will suppose henceforth that $\phi_i^1(t) - \phi_i^2(t) \geq 0$, which means that a forward move is always profitable to the Seller. This constraint is stable under all fusions considered so far.

Likewise: $\beta_i^1(t) \geq 0$ is the active payment i needs to make to i 's Buyer to cancel the deal; and $\beta_i^2(t)$ is the (possibly negative) passive payment from i 's Buyer to i which follows. When i plays that backward move, i 's payoff is therefore: $\beta_i^2(t) - \beta_i^1(t)$.

For a position $v \in \gamma$ we define upstream looking bounds:

$$\begin{aligned} \beta(v, t) &= \max_{i < v} (\sup_{s \geq t} (\beta_i^2(s) - \beta_i^1(s))) && \text{passive expense on a } B\text{-move} \\ \phi(v, t) &= \sum_{i < v} \sup_{s \geq t} (\phi_i^1(s) - \phi_i^2(s)) && \text{active expense on an } F\text{-move} \end{aligned}$$

where maxes and sums leave out positions for which the quantities of interest are undefined - eg an $i < v$ such that $|\phi_i| \cap [t, +\infty) = \emptyset$ indicating an arc that cannot be in an F -state.

The idea is that $\beta(v, t)$ is an upper bound for the expenses v may incur upon eviction by a B -move, whichever is the trace followed. Likewise, $\phi(v, t)$ is an upper bound for the price v will ever have to pay to acquire the underlying (by buying all positions upstream using a series of F -moves). Both quantities depend only on the arcs upstream of v in γ . (Incidentally, as there are no rules for inserting arcs other than at the end, the set of arcs upstream of v never increases as the tradeline evolves or time advances.)

We define also downstream looking bounds:

$$\eta_v^B(t) = \sup_{s \geq t} (\beta_v^1(s) - \beta_v^2(s)) \quad \text{active expense on a } B\text{-move}$$

The idea is that $\eta_v^B(t)$ upper bounds the active payment made by v on a B -move. This bound will work also for the semi-ring backward fusion, as long as player is not fool enough to pick an option worse than his original one. Here the control is local to v , ie does not depend at all on the trade line, because v 's β clause propagates via the standard μ_{BB} when triggered (see 8). Note that $\eta_v^B(t) \leq 0$ if v 's backward clause always specifies a profit for v (as in the standard bilateral contracts). There is no need to define a symmetric $\eta_v^F(t)$ to control for passive expenses on an F -move, as we have assumed above that F -moves are always profitable to the Seller.

Proposition 2 (max expenses) *Let γ be a trade line, and let v be a position in γ . Along any standard trace starting from γ where v extends the trade line (meaning plays no S_{vw}), v 's expenses are upper bounded by:*

$$\phi(v, t) + \beta(v, t) + N\eta_v^B(t)$$

with N the number of B_v moves played by v in the trace. In particular, if $\eta_v^B(t) \leq 0$, the expenses of v are upper bounded by $\beta(v, t) + \phi(v, t)$.

To cope with general traces where it may also happen that v extends the trade line, ie plays possibly multiple moves of type $S_{vw}(\beta_k, \phi_k)$, we can readily modify the estimate $\eta_v^B(t)$ to also maximise over such moves k :

$$\hat{\eta}_v^B(t) = \max_k \sup_{s \geq t} (\beta_k^1(s) - \beta_k^2(s))$$

If all these payments are Seller-positive, ie $\hat{\eta}_v^B(t) \leq 0$, we can forget this term as we did in the statement above.

The fact that expenses are unbounded if v 's Seller contract specifies that v has to pay to trigger a backward move (so the opposite of a penalty) should not overly concern us (although maybe it should concern v 's owner!). Indeed this cost is under the control of the player owning v . The same remark applies to the expenses incurred by extending repeatedly the end of the trade line.

The above statement gives a direct upper bound on the prepayments needed for a player to join the game. Depending on the specific time-dependencies - some of this provision can be returned as time advances and provisions are re-evaluated. Predictability is convenient here, as one knows the amount to provision.

The same inductive proof (omitted) also shows the following reasonable property, namely that the evolution of a trade line cannot lead to a solution where a Seller receives less than the originally asked price for a forward as well as a backward move.

Proposition 3 (Monotonicity) *Let γ be a trade line, and let v be a position in γ : v 's forward payoff (as Seller) can only increase, and the backward payment made by v (as Buyer) stays invariant.*

Regarding the first statement, one should not construe it as saying that the forward payoff will necessarily happen. Even if there are Buyers waiting to join, that is. To see this, suppose v faces a w who extends the tradeline with a forward-dead contract $S_{wx}(\mathbf{0}, \mathbf{1})$, and w never plays F_w , the only way out for v - a way which we know exists by general results explained next Section - is to B -sweep the trade line entirely (using additive forward fusions, at least).

The situation with the second statement is different as there is the B -sweep guarantee, namely that the entire trade line can be recalled by a position v under v 's original condition β_v . However, as we have seen, the sweep may have a price if $\eta_v^B > 0$.

We can also remark that the number of F -steps for v to obtain control of the underlying is upper bounded by the length of γ at the time v enters (necessarily as a Buyer); and, depending on the fusions used, it may become cheaper (for fixed time) as players upstream move forward.

Thus players can upper bound their costs. This is akin to so-called 'safety properties', namely that something bad never happens. Now we turn to the question of 'liveness', ie whether something good can happen.


6 Confluence of the trade game

Clearly, starting from a given trade line, and even assuming no-one else enters the game, players are able to resolve the tradeline in fundamentally different ways. What happens if one freezes time as well as extension moves? This 'frozen game' still has exponentially many moves (in the length of the trade line). Perhaps surprisingly, we will now show that, if the game is BF -exclusive, there is at each instant in time a unique maximally contracted form of the trade line which is reachable by the players. There is no guarantee that this irreducible form is a resolution of the trade line, neither is the result saying that the players will want to reach this normal form. Nevertheless, it is an interesting structural property of this class of games: the local property of BF -exclusivity generates a global property of stability.

To establish this we study *critical pairs* in the game, that is to say pairs of contractions (of the same type or not) which are both concurrently possible and can interact. For each of these, we exhibit a pair of contractions which converge back to the same trade line. Because of the dependency of the individual payoffs in the actual path, players may have conflicting preferences for the paths to be followed.

► We proceed axiomatically and use generic fusions for which we derive sufficient conditions to prove the above results.

Backward critical pair B_1B_2 We start by examining the case where both moves involved in the critical pair are of the same type. Let us look at the backward case first. Two B -moves interact non trivially only if their triggering positions u_1, u_2 are contiguous. We call this first critical pair B_1B_2 in reference to the two positions triggering the concurrent moves (indicated in red below). Transitions in the diagram are labelled by the type of move and the clause being triggered (time not shown):



For both legs of the opening span B_{u_1}, B_{u_2} to exist at the same time, and therefore constitute an actual critical pair, it must be that both β_1 and β_2 hold in the current context. One can close this span of transitions with the co-span B'_{u_1}, B_{u_1} provided $\mu_{BB}(\beta_1, \beta_2)$ also holds. This *co-timeliness condition* can be written $\beta_1 \wedge \beta_2 \Rightarrow \mu_{BB}(\beta_1, \beta_2)$ -where the implication is simply the inclusion of time domains.

To obtain confluence on the state, we also need both paths to lead to the same forward and backward clauses. Clearly, this amounts to asking that the backward fusion operators μ_{BB}, μ_{BF} are associative.


Payoffs combine passive and active payments and depend on the paths followed. To record the difference between two paths we introduce the *payoff commutator*:

$$B'_{u_1}B_{u_1} - B_{u_1}B_{u_2} = \mu_{BB}(\beta_1, \beta_2)u_1u_3 - \beta_2u_2u_3$$

With the projective variant of μ_{BB} , the commutator simplifies to $\beta_1u_1u_3 - \beta_2u_2u_3$. If, as in the standard contract, a backward effect is a penalty paid by Buyer, then u_2 prefers the right path where he receives this payment, while u_1 prefers the left path for the same reason, and u_3 prefers the path where he pays the least penalty (but he does not have a say).

In the case where u_4 does not exist and u_3 is the end of the trade line, the above diagram still makes sense and converges on u_1 (with no conditions), which becomes the new end.

Forward critical pair F_2F_3 We have a symmetric diagram for the (F_2F_3) critical pair:



The co-timeliness condition for the closing co-span is now $\phi_1 \wedge \phi_2 \Rightarrow \mu_{FF}(\phi_1, \phi_2)$. State confluence amounts to associativity of the forward fusions μ_{FB}, μ_{FF} . The payoff commutator is:

$$F_{u_3}F_{u_2} - F'_{u_3}F_{u_3} = \phi_1 u_1 u_2 + \mu_{FF}(\phi_1, \phi_2) u_2 u_3 - (\phi_1 u_1 u_3 + \phi_2 u_2 u_3)$$

It simplifies to $\phi_1 u_1 u_2 - \phi_1 u_1 u_3$ for the projective $\mu_{FF}(\phi_1, \phi_2) = \phi_2$. Prior remarks on the players' preferences apply *mutatis mutandis*. If forward clauses are straight payments, it is advantageous to u_3 to go down the $F_{u_3}F_{u_2}$ path, and let u_2 move first, as u_3 shares the forward price with u_2 . On the other hand, if the game uses the additive $\mu_{FF}(\phi_1, \phi_2) = \phi_1 + \phi_2$, the commutator contains now an additional contribution (in red) $\phi_1 u_1 u_2 + \phi_1 u_2 u_3 - \phi_1 u_1 u_3$. Thus, this variant makes paths indifferent to the players provided $\phi_1 u_1 u_2 + \phi_1 u_2 u_3 = \phi_1 u_1 u_3$, a reasonable quotient by transitivity of payments.

The diagram still makes sense if u_4 is absent and u_3 is the end.

We assume henceforth that our fusions are associative.


Tug-of-war critical pair (B_2F_3) There are three critical pairs mixing both types of contractions. The most obvious one is the ‘‘tug-of-war’’ critical pair (B_2F_3) where players contract the same trade arc. If both clauses ϕ_1, β_1 are indeed simultaneously fireable, there is no general way to close this critical pair.

$$u_1 \xrightarrow[\beta_1]{\phi_1} u_2$$

Imagine u_1 is the origin, and u_2 the end, then both moves lead to distinct irreducible forms.

► Henceforth we set the game in the BF -exclusive fragment and hence forbid the (B_2F_3) pair.

Self-critical pair B_2F_2 The next critical pair involves just one player:



The associated co-timeliness condition for closure is $\beta_2 \Rightarrow \mu_{FB}(\beta_1, \beta_2)$.


State confluence has now a variety of solutions, investigated below jointly with the next and last critical pair.

The payoff commutator is:


$$B_{u_2}F_{u_2} - F_{u_2}B_{u_2} = \mu_{FB}(\beta_1, \beta_2) u_2 u_3 - \beta_2 u_2 u_3$$

which is zero if $\mu_{FB}(\beta_1, \beta_2) = \beta_2$. This is natural. In other cases, there is potentially order-sensitiveness in that the block-maker will decide the payoff of u_2 . However, u_2 is directly responsible for signing on both trade arcs, so may mitigate this dependency at the strategic level.

Brokering critical pair B_1F_3 The last critical pair (B_1F_3) is the most interesting, in that it is not a simple commutation, and one has two closing options, either the B co-span (in green):



or the F one (in red):



The associated co-timeliness conditions for closure depends on which closing option one chooses:

$$\begin{aligned} \beta_1 &\Rightarrow \mu_{BB}(\beta_1, \beta_2) && B \text{ co-span} \\ \phi_2 &\Rightarrow \phi_1 \wedge \mu_{BF}(\phi_1, \phi_2) && F \text{ co-span} \end{aligned}$$

Note that the first condition implies the one obtained earlier for the B_1B_2 critical pair. As to the F co-span, the co-timeliness constraint couples ϕ_2 to ϕ_1 . Worse, under the BF -ex condition, it cannot be that ϕ_1 holds, since β_1 does, so the F co-span option to close is not available. For this reason, we continue our analysis with the B co-span only.

The payoff commutator is:

$$B_{u_1}F_{u_3} - B'_{u_1}B_{u_1} = \beta_1 u_1 u_2 + \mu_{BB}(\beta_1, \beta_2) u_1 u_3 - (\phi_2 u_2 u_3 + \beta_1 u_1 u_3)$$

which simplifies to $\beta_1 u_1 u_2 - \phi_2 u_2 u_3$ for the projective variant. Anticipating on the strategic analysis of the standard ternary game given in §7, we can see that the B_1F_3 critical pair embeds in this game, assuming u_3 is the end of the trade line. Specifically, $F'_{u_3}F_{u_3}$ is the green branch in Fig. 3, while $F_{u_3}B_{u_1}$ is the red one (the one to be avoided if the game is to be of interest to a broker). In the light of this embedding, we can map this commutator to the special case where $\beta_1 u_1 u_2$ is the standard penalty owed by u_2 , and $\phi_2 u_2 u_3$ is a commission paid by the final Buyer. In this case, the commutator has a strategic interpretation, in that it measures, from the point of view of the middle player, the difference between a satisfactory brokering process where the player is paid the expected commission, and one where he is not.

Solving for confluence With all four critical pairs in place, we see that solving confluence leads to the same pair of equations:

$$f(g(x, y), z) = g(x, f(y, z)) \quad (9)$$

$$f(f(x, y), z) = f(x, g(y, z)) \quad (10)$$

with $f = \mu_{BF}$, $g = \mu_{FF}$ for the forward fusions, and $f = \mu_{BB}$, $g = \mu_{FB}$ for the backward ones.

Thus, so far as confluence is concerned, one can choose independently the forward and backward fusions. Moreover, any symmetric assignment $\mu_{FF} = \mu_{BF}$, $\mu_{FB} = \mu_{BB}$ is a solution as long as the fusions chosen are associative and satisfy the co-timeliness conditions. For instance fusions based on max may lead to interesting dynamics. If we restrict to linear fusions, we can characterise the non-degenerate solutions to confluence:

Proposition 4 (confluent fusions) *Associative and linear confluent fusions are of the following forms:*

$$\begin{aligned} f(x, y) = g(x, y) &= x + y \\ f(x, y) = x, g(x, y) &= cx + dy \\ f(x, y) = g(x, y) &= y \end{aligned}$$

with $c, d \in \{0, 1\}$ and $c + d > 0$; $f = \mu_{BF}$, $g = \mu_{FF}$ for the fwd fusions; $f = \mu_{BB}$, $g = \mu_{FB}$ for the bwd ones, with in this case only the first two forms available and with the additional restriction that $d > 0$. Moreover the combination $\mu_{FF}(x, y) = x$, $\mu_{FB}(x, y) = y$ is forbidden.

Among the linear solutions, we recognise the standard fusions $\mu_{FF}(x, y) = \mu_{BF}(x, y) = x + y$ (additive fwd fusion), and $\mu_{BB}(x, y) = x$, $\mu_{FB}(x, y) = y$ (projective bwd fusion).

6.1 Fixed-time confluence

It is easy to verify that we have already considered all critical pairs: the three mixed ones: B_1F_2 , B_1F_3 , and B_2F_2 , and the two pairs of the same type: B_1B_2 and F_2F_3 . Critical pairs involving the S -moves are trivially closable. Other pairs of concurrently fireable moves trivially commute as they have no overlap. It follows that the dynamics of the fixed-time game is locally confluent in the BF -exclusive fragment (which is closed under all moves), which forbids the bad pair B_2F_3 . As in addition spans are closed by co-spans of length one, contractions form a confluent system at fixed time (up to payoffs). In particular:

Proposition 5 (fixed-time confluence) *Let γ be a BF -exclusive tradeline, there is a unique irreducible form γ_0 reachable from γ in the fixed-time closed game.*

Thus maximal plays at fixed-time (or during any lapse of time where all clauses in the game are constant) lead to exactly the same trade line in the closed game. Besides, as this restricted game is clearly finite, there is indeed a unique reachable irreducible form. When time clicks, however, γ_0 may cease to be irreducible. Our machinery for reversible and composable trades is all about timing. Indeed: (i) the one potential winner will change over time in general; (ii) not because there is a unique normal form means that the players want to reach it -some will want to run the clock for better payoffs; finally, (iii) not because there is a unique normal form means players can reach it -to exploit the confluence property their pockets must be deep enough. Another way to phrase (iii) is to say that the local confluence analysis is independent of the effects/payoffs, and only depends on the logical part of clauses (to make sure pairs are closable in the same instant).

6.2 Game-theoretic confluence

Looking back at the critical pairs (CPs), we see that they contain more information than local confluence: in each case, there is a prevailing player which can perform a move regardless of whether his initial move is first past the block-maker post or not. We capture this idea of game-theoretic confluence in a stronger statement of confluence: at fixed time, there is a unique ‘focal’ player who can, under additional assumptions, drive the game all alone to resolution. (The usual caveat applies: not because the focal player can collapse the trade line means that he wants to.)

Recall that we say a trade line γ is connected in a given context if none of its arcs is idle. If γ is also exclusive, this means that every arc in γ has a definite orientation.

Lemma 2 *Connectedness is preserved by contraction under standard fusions (see Def. 8).*

Definition 10 (focal position) *Let γ, γ' be trade lines. We write $\gamma \Rightarrow^* \gamma'$ for a sequence of moves (possibly empty), ie a trace, leading from γ to γ' . Let u be a position in γ , t a time. Say a trace $\gamma \Rightarrow^* \gamma'$ is a u -trace at t if it consists entirely of u -moves executed a time t . Say u is a focal point of γ at t if there is a (maximal) u -trace at t which resolves γ to position u .*

A focal position is one (the owner of which) can collapse the trade line all alone. When time clicks, it may change, but at a given time it is unique (by confluence).

Lemma 3 (lifting) *Suppose γ is BF-exclusive and $\gamma \Rightarrow^{u^\epsilon} \gamma'$ (with ϵ the move’s direction). Then 1) u is focal in γ at t iff u is focal in γ' at t ; 2) if u is focal in γ at t then all maximal u -traces lead to u .*

Notice that: 1) the above statement makes sense as u exists in γ' because u^ϵ does not evict u from the game; 2) is not trivial in that we know that all maximal traces lead to the same normal form (by confluence); but some maximal u -traces may not be maximal as traces.

Proposition 6 (pacman) *Let γ be a BF-exclusive trade line. If γ is connected at time t , it has a focal position at t . Furthermore, if $\gamma \Rightarrow^* \gamma'$ and u is focal in γ , then u is focal in γ' .*

Hence, not only is there always a focal player who can collapse the tradeline on his position at fixed time - but, *whatever the other players do*, the focal player retain that power. Concretely, it is easy to see that the focal position in a connected trade line is the position furthest away from the origin upstream of which all arcs are in an F -state.

7 The standard ternary game

We now consider two chained copies of the standard bilateral contract (§6) with a view to understanding the strategic aspects of this particular game, ie how it is played by rational players. We suppose each copy has its own independent parameters. The direct way to build this ternary game on positions u, v, w is to follow the sequence $S_{uv}(\Delta, p; a); S_{vw}(\Delta', p'; a')$:

$$\cdot u \cdot \xrightarrow{S_{uv}} \cdot u \xrightarrow[\Delta, p]{a} v \cdot \xrightarrow{S_{vw}} \cdot u \xrightarrow[\Delta, p]{a} v \xrightarrow[\Delta', p']{a'} w \cdot$$

Payoffs include a ‘pay-per-move’ amount ϵ which abstracts running costs (gas). Costs associated to setting up the game are not considered. We write \odot for the asset underlying the trade line. We also write $\mathcal{U}_x(\alpha \odot + a)$ for the utility function of the player at position x with $\alpha \in \{0, 1\}$ representing an amount of the

underlying asset, and a some amount of cash (possibly negative). We suppose \mathcal{U}_x is non-decreasing in both arguments α and a .

Our specific question is whether the game provides liquidity. That is to say, is it possible to set its parameters in such a way that it is profitable for a pure broker (a player who will not buy the asset) to enter the game at the middle position v assuming rational co-players. We leave to further research the analysis of the game under incomplete information about players' preferences.

Game analysis For the overall trade to be Pareto-improving, players must have different preferences for the asset. In the language of preferences (or utilities), this assumption reads:

$$\begin{aligned} \mathcal{U}_u(\odot) &< \mathcal{U}_u(a) \\ \mathcal{U}_w(\odot) &> \mathcal{U}_w(a + a' + 2\epsilon) \end{aligned}$$

The second inequality says that w is willing to play the sequence $F_w; F_w$ and thus to fully resolve the trade line and acquire the asset. The first one says that u is willing to sell the asset at the price offered to v .

The initial ternary trade line has 8 resolutions. Let us start with the resolutions where v does not play at all. The corresponding sub-game is shown in Fig. 3 with the payoff vectors indicated at the leaves. There are two plays where w gets the asset: a full forward play $F_w; F_w$ (in green), or a play $B_u; F_w$ (in red) where the mediating player v is evicted by u . There are two plays where u cancels the trade line and recovers the asset: a full backward play $B_u; B_u$, and a play $F_w; B_u$ where w is frustrated in her attempt to buy the asset. The opening span is an instance of the $B_1 F_3$ critical pair (§6).




Figure 3: Game tree with passive broker: Δ is the time at which the B_u cancellation option of u becomes active (indicated by the conditioning notation $B_u | \Delta$). The rightmost (green) branch is the desired execution of the trade line. The leftmost one (red), which for certain parameters is the subgame perfect equilibrium, one would rather avoid as it robs the broker from his fee.

Backward moves are only available at Δ , Δ' and onwards. Backward induction shows that the (red) sequence $B_u; F_w$ is the only subgame-perfect equilibrium as long as $a' + p > \epsilon$. This condition is mild since it assumes that computational costs ϵ are lower than the sum of the broker's fee a' and punishment p for not closing the deal. Obviously, this is not the intended equilibrium since the broker v is bypassed (and even punished for not forwarding the trade). Intuitively, the buyer w does not exercise her forward option before Δ because she knows that the seller u will close the deal on his own, by first triggering his backward option and then letting w execute her forward option to buy the token. The benefits for u is that he gets the broker fee (and the punishment), while w does not have to pay the computational costs of the first move which are now covered by u .




Figure 4: Game tree of broker: Δ' is the time at which v 's cancellation option becomes active. The blue branch is a deterrence path, not played, yet useful to incite w to play the other subgame.

This result indicates that the standard ternary game is not incentive compatible when the middle player v is passive. Consider now the game tree when the broker v moves first. Focus on the leftmost (blue) branch of the tree in Fig. 4. From Δ' onwards, v can exercise her option to cancel the trade with w . (Crucially, without reversibility, v has no such option!) This move guarantees that w receives a negative payoff, and is therefore a threat. Setting the penalty $p' > p + \epsilon$ ensures that v will always get a positive payoff, which makes the threat credible. Hence a rational w will never wait past Δ' as v would exercise her backward option and punish w for trying to bypass her intermediation. The cancellation option of v deters w from playing the “bad” subgame perfect equilibrium depicted in Fig. 3. Note, however, that deterrence is not operational if u can expel v before v can expel w , that is to say if $\Delta \leq \Delta'$.

Proposition 7 *The desired execution $F_w; F_v$ is the only subgame perfect equilibrium of the standard ternary game when (i) $\Delta > \Delta'$, and (ii) $p' > p + \epsilon$.*

By the same reasoning, one sees that v can resist a collusion between u and w in which they try to extort p from v , with w having no intention of buying. As v can always do B_v after Δ' and stop there with a guaranteed positive payoff. The B_v deterrence move works its magic: v will not play that move, but the fact that she *may*, drives the equilibrium where she wants it. Importantly, the suitable constraints on parameters can be met at the time of the extension move S_{vw} , ie when Δ', p' are chosen by v and agreed by w , because, then, Δ and p are known. Even more importantly, v does not need to have the cash to buy the asset, ie *zero-collateral intermediation* is feasible. It is also interesting to see that the inequalities are purely structural, ie they do not depend on the value of the underlying.

It is *only* because the standard bilateral contract is extensible that more players (vs) may want to join the trade line than those willing to enter the single game. The analysis above shows that a v with just enough cash to provision for the game will join the game.

The two inequalities in Prop. 7 are necessary since the game is not incentive compatible when they are violated (meaning rational players will not follow the desired execution). Whether they are also sufficient conditions with less stringent assumptions is a more complicated matter. In particular, it seems important to also consider situations where w 's willingness to pay is uncertain. Then one would like to prove that w cannot gain even if she is of the ‘impostor’ type, ie if she does not prefer the asset over $a + a'$, so that $U_w(\odot) < U_w(a + a' + 2\epsilon)$. To discuss the problem informally, notice that at Δ' , w loses exclusivity, as v can at any time onwards fire B_v , and then resample the pool of possible buyers by reversibly selling its position to find a more reactive Buyer. Reversibility allows v to repeatedly kick out slow ws by playing

B_v at Δ' and obtain payoff (with timeouts expressed as delays):

$$\pi_v \geq -p + \lambda(p' - \epsilon)(\Delta'/\Delta) \geq -p$$

where $\lambda < 1$ (which depends on the spread a') measures the frequency at which v resamples, ie how good a broker v is for a given level of demand, and $-p$ is a lower bound on v 's loss (in the case v finds no w at all during Δ). If v never finds a true Buyer during the allotted time Δ , despite resampling as soon as possible, v will still break even provided $\Delta p \leq \lambda \Delta' (p' - \epsilon)$.

8 Aside on implementation

In a direct implementation of our protocol, the trade line and its evolution rules are interpreted by a dedicated smart contract, relying on an external custodial contract to define the ownership of the assets used in the trade game. These assets can be security tokens or currencies in the popular ERC20 model from Ethereum. When the game starts, the owner of the asset transfers its property to an account of the custodial contract which is controlled by the interpreter contract. When the trade line finally resolves, it remains for the interpreter contract to ask the custodial contract to transfer the ownership of the asset to the owner of the one remaining position in the game (which implies that the owner is known to the custodial contract). To implement passive payments, one can forward payment obligations to an external system managing the players' debts. Another more decentralised approach is to ask players to pre-provision for passive payments at the time of joining the game, so that there is no risk of them not being paid. The max expense Lemma (§5) shows this can be done. There is no such concern for active payments, as these are payments which players have to make to change the state of the game. Players will do as they please.

9 Conclusion

We have defined a trade protocol to manage chains of reversible bilateral contracts. Its design derives entirely from a simple premise: the need for a theory of deferred payments which allows one to postpone payments, and resell an asset one has not paid for yet. To do this there is no magic, one has to keep somehow a memory of past transactions, and be able, as the need may occur, to revert some. This leads us to a protocol where each chain of transactions, or trade line as we call them, defines the state of an open game; the evolution of which relies on the reversibility of the component games. Our framework is as a consistent calculus of deferred payments and novation. Contracts which one can express therein are expressive enough so that novel financial objects can be built and experimented with, and yet simple enough that contracts are financially sound and have a transparent cognitive model.

There are many avenues to enrich and refine our trade games. One can explore other forms of contractions for the trade line. For example, it may be that one can emulate directly auction structures by using max-based clause fusions instead of linear fusions. Another change in the rules of the trade game would be to introduce batched extensions such as $S_{uw}; S_{vw}$ where w and u are positions held by the same player. These are already possible over-the-counter and can improve fungibilisation of positions in a trade line. Another more speculative direction is that of splitting trade lines. The intention is to let players enter a given trade line competitively. Competing here amounts to allow for non-exclusive extensions of the trade structure. This leads to *trade trees* and can perhaps further improve liquidity. It remains to be seen whether one can set up such an extension in a sound way.

References

- [1] Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [2] Drew Fudenberg and Jean Tirole. Game theory, 1991. *Cambridge, Massachusetts*, 393(12):80, 1991.
- [3] Mariana Khapko and Marius Zoican. How fast should trades settle? *Society for Financial Studies (SFS) Cavalcade*, 2017.