

# Engineering Foreign Exchange Processes via Commitment Protocols\*

Nirmit Desai<sup>§</sup>, Amit K. Chopra<sup>§</sup>, Matthew Arrott<sup>¶†</sup>, Bill Specht<sup>‡‡</sup> and Munindar P. Singh<sup>§</sup>

<sup>§</sup>North Carolina State University, Raleigh, NC 27695-8206, {nvdesai, akchopra, singh}@ncsu.edu

<sup>¶</sup>University of California at San Diego, La Jolla, CA 92093, marrott@ucsd.edu

<sup>‡</sup>Currenex, Inc., 1700 Seaport Blvd., Redwood City, CA 94063, bill.specht@currenex.com

## Abstract

*Foreign exchange (FX) markets see a transaction volume of over \$2 trillion per day. A number of standard ways of conducting business have been developed in the FX industry. However, current FX specifications are informal and their business semantics unclear. The resulting implementations tend to be complex and compliance with the standards unverifiable. This results in potential loss of value due to incompatible business processes and possible trades not consummated.*

*This paper validates a formal, protocol-based approach by specifying foreign exchange processes as standardized by the TWIST consortium. The proposed approach formalizes a small, core set of foreign exchange interaction protocols on which the desired processes can be based. The core protocols can be composed to yield a large variety of possible processes. Each protocol is rigorously defined in terms of the commitments undertaken and manipulated by the parties involved. By contrast, traditional approaches as used in the current TWIST specification lead to redundancy in specification and difficulty in understanding the import of the interactions involved. In addition, our approach discovered interesting business scenarios that traditional approaches would have missed.*

## 1 Introduction

With *daily* total traded cash volume to the tune of \$2.3 trillion, FX markets are huge [20] (over double the US stock market). Owing to the growth in international business and the globalization of enterprises, FX trading has increased by more than 30% in the last year and has more than doubled since 2001 [3]. Electronic trades—a major factor behind

this growth—are expected to account for more than 75% of all global FX trades by the year 2010 [2].

To support such growth in electronic trading, it is crucial to standardize key messages and workflows. TWIST (Transaction Workflow Innovation Standards Team) Process Innovations is a not-for-profit industry group of corporate treasurers, fund managers, banks, system suppliers, electronic trading platform and market infrastructure vendors, and professional services firms. TWIST has collaborated with industry partners and standards organizations such as FPL (FIX Protocol Ltd.) [9] and ISDA (International Swaps and Derivatives Association) [12] to define standard good practice processes throughout the transaction processing life cycle for wholesale FX trades [19].

Current FX standards specifications describe business processes informally in the form of natural language descriptions accompanied by sequence diagrams representing typical scenarios. FX processes (as in other businesses) typically have several dimensions of variation, e.g., trading with or without credit checks and trading with or without a trading service. Owing to their informal nature, existing FX specifications treat each scenario as a separate case despite commonalities, making it harder to determine the relationships among such variations and whether they can be combined to serve a particular need. As a result, a large number of processes are explicitly specified. Managing and—equally importantly—understanding such large sets of standards is difficult.

Interaction-oriented approaches represent a growing trend in business process modeling [17, 21, 10, 23]. RosettaNet's Partner Interface Processes (PIPs) support billions of dollars of business each year [15]. PIPs are interaction-oriented, but informally specified and limited to two-party request-response interactions.

By contrast, this paper advocates a formal approach to modeling interactions [6, 8]. Interactions among parties are treated as first-class modeling abstractions. The formal semantics enables the reuse, refinement, and composition of the core interaction patterns. To emphasize the contrac-

\*With partial support from the US National Science Foundation under grant IIS-0139037.

<sup>†</sup>Matthew Arrott is TWIST CTO.

<sup>‡</sup>Bill Specht is TWIST Wholesale Trade Prime.

tual semantics involved in such processes, commitments among the parties are explicitly modeled. The messages are formalized in terms of how they affect the commitments among the partners.

The goal in this paper is to identify a set of core interaction patterns and formalize them as protocols. Such protocols can be composed with each other in a variety of ways to derive the large set of possible combinations. Thus, a large set of processes can be engineered using a small set of modular protocols. More importantly, new business scenarios are discovered while composing protocols. This exercise helps identify ambiguities and gaps in the specification. Given the sheer scale, variety, and critical nature of FX transactions, the impact of such an engineering approach can be enormous.

This paper (1) validates the protocols approach [7] via an extensive knowledge engineering exercise, and (2) develops a methodology for creating protocols. We show that such informal standards specifications (TWIST and others) are ambiguous and incomplete. Moreover, they do not adequately support modularity and composition—we found that 28 TWIST processes can be specified in terms of 12 core protocols. Moreover, this 12 core protocols can be used to specify novel processes not described in the specification.

Section 2 outlines TWIST processes for price discovery. Section 3 models the commitments in price discovery and points out some of its shortcomings. Section 4 shows how the TWIST processes may be obtained by composing elementary protocols.

## 2 Price Discovery Processes

This section describes the price discovery processes from the TWIST specification [19, Sec. 7.2]. Figure 1 describes bilateral price discovery processes 7.2.1 and 7.2.2. A *Taker* is trying to discover the price of a currency in another currency; the *taker Maker* provides the price. The *Maker* indicates in the *priceResponse* if an execution confirmation is required. Here, *executionConfirmation* means the quoted and accepted prices are agreed upon and the deal is reached. If confirmation is required (7.2.2), the quoted price is not binding to the *Maker* even if the *Taker* accepts it. Otherwise (7.2.1), the *Maker* is bound to trade at the quoted price if the *Taker* accepts it. In either case, the *Taker* may not reveal whether the requested currency is to be bought or sold, forcing the *Taker* to respond with both bid and offer quotes, thereby revealing the spread. The *Taker* can choose the direction of the trade when he accepts any of the quotes.

In addition to the messages shown in the figures, the *Taker* may cancel a request by sending *cancelPriceRequest*, or reject quotes by sending *nothingDone*. Also, the *Maker* may cancel a quoted price by sending *cancelPrice*. A *priceResponse* may also indicate a refusal to quote a price,

and an *executionConfirmation* may also indicate failure to execute the deal.

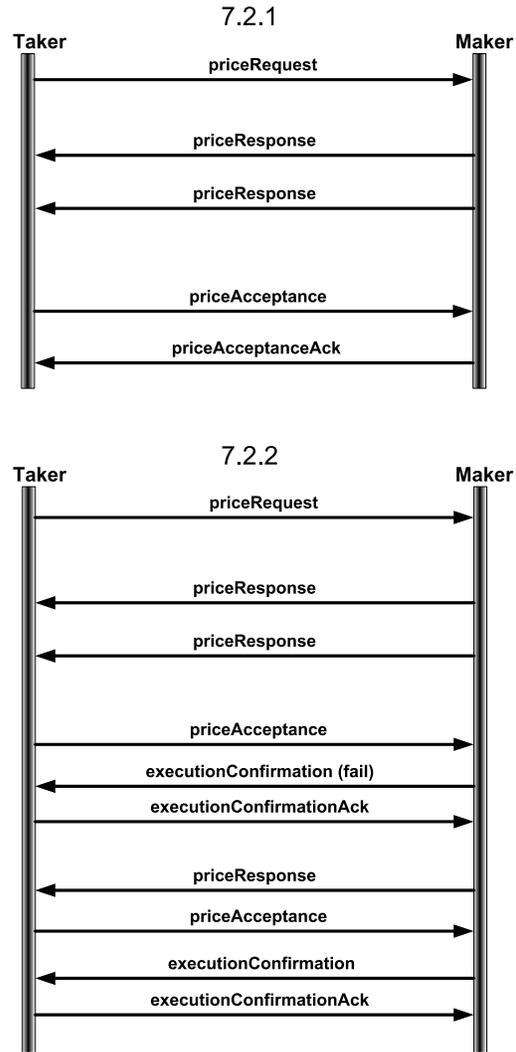


Figure 1. Bilateral RFQ (7.2.1 and 7.2.2)

Figure 2 describes multilateral price discovery wherein the *Taker* uses the trading service to discover the best price. The trading service may interact with multiple *Makers* to find a price for the *Taker*. To the *Makers*, this scenario is identical to the bilateral case. To the *Taker*, the only difference is that it receives responses from multiple *Makers*.

## 3 Protocols for Price Discovery

This section discusses the limitations of semiformal specifications (such as the current TWIST documents [19]) in terms of semantics, verifiability, and precision. It then formally specifies TWIST processes 7.2.1 and 7.2.2.

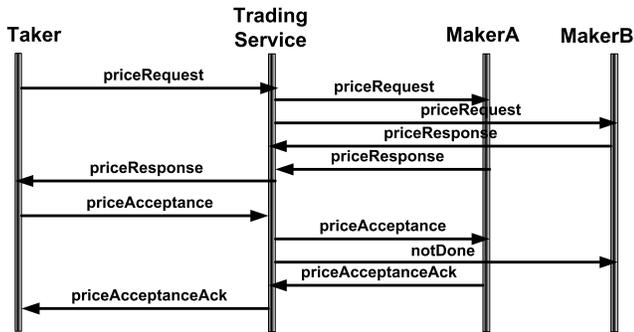


Figure 2. Multilateral RFQ (7.2.3)

### 3.1 Commitments in Price Discovery

A commitment  $cc(x, y, p, q)$  obligates a *debtor*  $x$  to a *creditor*  $y$  for fulfilling the condition  $q$  if  $p$  holds [16]. Here  $p$  is the *precondition* and  $q$  is the *condition* of the commitment. When the precondition is true, the commitment is termed a base commitment, else a conditional commitment.

Commitments can be manipulated: create, discharge, to-Base (change to a base commitment), delegate (changing the debtor), assign (changing the creditor), release (creditor releasing the debtor from the commitment), and cancel (debtor canceling the commitment). Consider, for example, a scenario where a buyer and a seller are exchanging goods for payment. A conditional commitment  $cc(\text{buyer}, \text{seller}, \text{goods}, \text{payment})$  denotes an obligation from the buyer to the seller that if the goods are delivered, the buyer will pay. In the event that the precondition goods holds, the conditional commitment changes to a base commitment  $cc(\text{buyer}, \text{seller}, \text{true}, \text{payment})$ . In the event that payment holds, the buyer's commitment is discharged and does not hold anymore. Commitments do not imply temporal ordering, e.g., payment may happen before goods, thus, discharging the commitment.

Protocols declaratively specify choreography of the messages exchanged among roles. They give messages a contractual semantics by defining how they affect the participants' commitments. For example, a message signifying shipment may cause the precondition goods, thereby causing the commitment to change to a base commitment. As a conversation progresses, commitments among the parties change to represent its evolving contractual state. Unless the precise meaning of the messages in terms of how they affect the extant commitments is specified, ambiguities may ensue about the participants' obligations.

Assuming that the *Taker* is selling currency  $\text{cur1}$  to the *Maker*, Figure 3 depicts various interpretations of the messages in process 7.2.1. The boxes denote states consisting of the commitments holding: these are newly created commitments and commitments from previous states (for brevity,

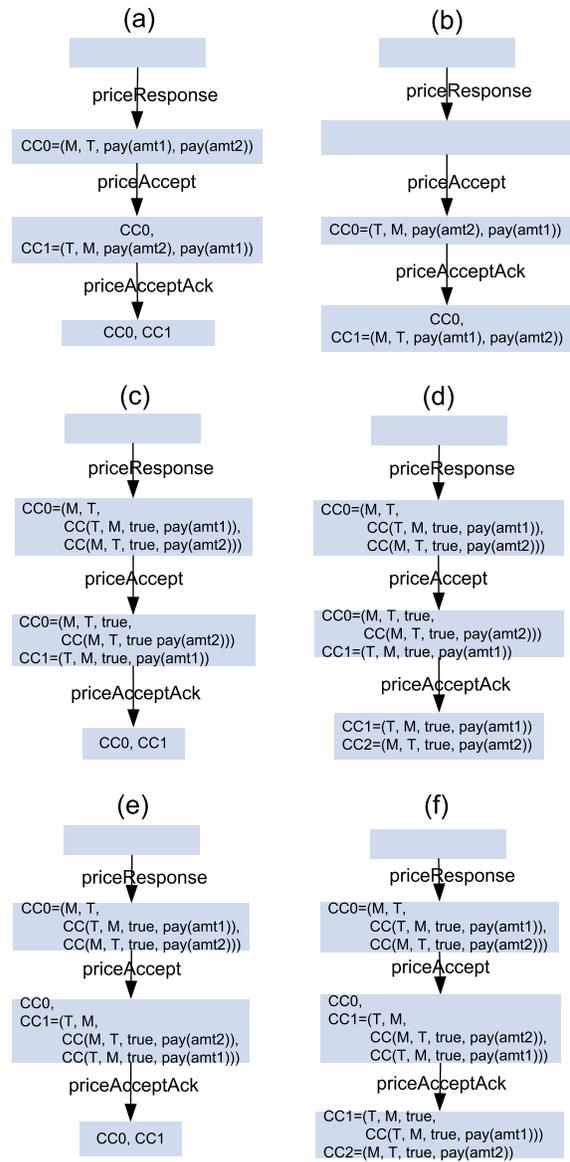


Figure 3. Possible interpretations of some TWIST messages in terms of commitments

only the new commitments are written out in detail). The *Maker* is denoted by  $M$  and the *Taker* by  $T$ . All the cases result in a state in which both parties have committed to each other for payment. However, there are subtle semantic differences in terms of how the risks and benefits of the participants evolve [22]. Also, nested commitments as in (c), (d), (e), and (f) allow more flexibility than the un-nested commitments as in (a) and (b) (as we shall see shortly).

In (a), the *Maker* and the *Taker* commit to each other to paying via  $\text{priceResponse}$  and  $\text{priceAcceptance}$ , respectively. However,  $\text{priceAcceptanceAck}$  is superfluous in the

sense that it does not affect the commitments. The final state is safe: regardless of the temporal ordering, payment is guaranteed if the commitments are not violated. In (b), `priceResponse` provides a nonbinding price. Whereas `priceAcceptance` has the same meaning as in (a), `priceAcceptanceAck` means that the *Maker* commits to the deal.

In (c), `priceResponse` creates a nested commitment: the *Maker* commits to the *Taker* to committing to paying if the *Taker* commits to paying. The condition and the preconditions being commitments enables their delegation, assignment, and so on. For instance, the parties need not make payments themselves but may delegate the commitments to their banks. Without nesting, such flexibility would be lost. Precondition `CC1` is caused by `priceAccept`, and `priceAcceptanceAck` is superfluous as in (a). The final state is safe as the *Taker* has committed and the *Maker* must commit to satisfying the nested commitment. In (d), the only difference from (c) is that `priceAcceptanceAck` is not superfluous and creates `CC2`.

In (e), the difference is that `priceAcceptance` causes a counter nested commitment instead of causing the precondition of `CC0`. Like in (c) and (a), `priceAcceptanceAck` is superfluous. In (f), the only difference is that the acknowledgment causes `CC2`—the condition of `CC0`.

The semantic differences among these variations highlight the importance of specifying such processes formally via commitments. Similar interpretations exist for TWIST process 7.2.2. A benefit of the present formalization exercise is that it helps identify the possible points of confusion and disagreement, which would otherwise have been glossed over in the documentation.

### 3.2 Gaps in Price Discovery

Although the typical scenarios are well-understood, an exhaustive set of possibilities can be covered rigorously only with formal methods. Our formalization exercise uncovered the following cases that are not clearly addressed in the specification. For example, does `nothingDone` terminate the conversation or merely reject a specific `priceResponse`? Also, can two price responses in a conversation have different confirmation requirements? Can the time until which the quote is valid be different for different quotes? Even more troublesome, what if `cancelPrice` and `priceAcceptance` cross? If the confirmation was not required, should the *Maker* be allowed to cancel a quoted price? Maybe in such a case, the *Maker* would use `priceAcceptanceAck` to indicate failure, but that is not obvious.

What if the *Maker* confirms execution but the *Taker* does not acknowledge it? Or what if the confirmation is not required and the *Taker* accepts a price but the *Maker* does not acknowledge it? The commitments should be specified in such a way that any unresponsive party is seen as violating

a commitment.

Answering such questions is critical; if they are not answered, then it reflects gaps in the specification. Our approach may not supply the answers to these questions, but helps uncover such questions via formal specification and verification techniques. This is the critical value of formal methods.

### 3.3 Specifying Protocols in $\mathcal{C}+$

Following Desai *et al.* [6], this paper specifies protocols in  $\mathcal{C}+$ , which is an action description language that gives primacy to causation [11].  $\mathcal{C}+$  supports *elaboration tolerance* enabling refinement of a specification merely by adding to the existing specification even if the desired effect is to disable some inferences. For protocols, elaboration tolerance means that certain interactions can be added, removed, or modified simply by adding axioms to an existing specification.

A  $\mathcal{C}+$  specification describes a transition system consisting of states and transition between them. A specification consists of a set of causal *laws*. A *fluent* is a proposition (true or false) whose value may change from state to state. *Actions* performed by agents *cause* the value of fluents to change. An inertial fluent continue to hold unless an action changes its value. The  $\mathcal{C}+$  semantics ensures that all and only the caused fluents hold at any state. The general concepts relating to protocols are specified in  $\mathcal{C}+$  as an ontology (Listing 1), to be included with specifications of individual protocols. The operator `++` denotes logical OR and `<>` denotes object inequality.

In  $\mathcal{C}+$ , an *exogenous* action is one that simply happens or not—the specification does not explain its cause. Messages are modeled as exogenous actions (line 9). The elaboration tolerance of  $\mathcal{C}+$  allows placing constraints on the order of action occurrences. Inertial fluents (line 8) record the effects of all message occurrences (line 26). A static fluent initial ensures that the start state of a protocol is void of any fluents or commitments (lines 15, 21–24). Static fluents are not inertial; in each state the value of static fluents is determined by the value of other fluents in that state.

Commitments are modeled as inertial fluents (line 12) and their preconditions and conditions are modeled as actions (line 10) that are disabled by default (line 29). Also, occurrences of conditions are recorded in inertial fluents (line 11 and line 28). For simplicity, Listing 1 only describes `create`, `discharge`, `cancel`, and `toBase` (lines 13–14). Whereas `discharge` and `toBase` are caused when the appropriate conditions hold, other operations are caused directly by the actions of the parties.

Causing the conditions and preconditions of a commitment causes appropriate operations: `discharge` and `toBase`, respectively, provided the commitment is active or being

created simultaneously (lines 31–35). If a commitment is discharged, it is deemed fulfilled and ceases to hold (line 37). If `toBase` is caused, the original commitment ceases to exist, and a base level commitment is created, only if the original commitment is not being discharged or canceled simultaneously (lines 42–45). A commitment is asserted if `create` is caused and that commitment is not being simultaneously discharged, converted to base, being canceled, and the commitment does not already exist (lines 47–48). All commitment operations are disabled by default (lines 50–53). These laws collectively ensure correct behavior of commitment operations in the face of concurrent actions.

### Listing 1. Protocol ontology

```

1 :- sorts Role, Slot; Message; Commitment; Condition.
3 :- variables
4   msg1 :: Message; p,q :: Condition;
5   ccl :: Commitment; dbl,cr1 :: Role.
7 :- constants
8   fl(Message) :: inertialFluent;
9   act(Message) :: exogenousAction;
10  cond(Condition) :: action;
11  fl_cond(Condition) :: inertialFluent;
12  comm(Commitment) :: inertialFluent;
13  create(Commitment), discharge(Commitment),
14  toBase(Commitment), cancel(Commitment) :: action;
15  initial :: sdFluent.
17 :- objects
18   T :: Condition;
19   CC(Role, Role, Condition, Condition) :: Commitment.
21 caused initial if initial.
22 caused -initial if comm(ccl).
23 caused -initial if fl(msg1).
24 caused -initial if fl_cond(p).
26 act(msg1) causes fl(msg1).
28 cond(p) causes fl_cond(p).
29 -cond(p) causes -cond(p).
31 caused discharge(CC(dbl, cr1, p, q)) if cond(q) &
32 (comm(CC(dbl, cr1, p, q)) ++ create(CC(dbl, cr1, p, q))).
34 caused toBase(CC(dbl, cr1, p, q)) if cond(p) & p <> T &
35 (comm(CC(dbl, cr1, p, q)) ++ create(CC(dbl, cr1, p, q))).
37 discharge(ccl) causes -comm(ccl).
39 cancel(ccl) & -discharge(ccl) causes -comm(ccl).
41 toBase(ccl) & -discharge(ccl)
42 & -cancel(ccl) causes -comm(ccl).
44 toBase(CC(dbl, cr1, p, q)) & -discharge(CC(dbl, cr1, p, q))
45 & -cancel(dbl, cr1, p, q) causes comm(CC(dbl, cr1, T, q)).
47 caused comm(ccl) if create(ccl) & -(discharge(ccl)
48 ++ toBase(ccl) ++ cancel(ccl)) & -comm(ccl).
50 -create(ccl) causes -create(ccl).
51 -toBase(ccl) causes -toBase(ccl).
52 -cancel(ccl) causes -cancel(ccl).
53 -discharge(ccl) causes -discharge(ccl).

```

Messages, as exogenous actions, can happen on any transition by default. Protocols typically specify a set of restric-

tions on such messages and rules for their effect on commitments. As there are commonalities in processes 7.2.1 and 7.2.2, a common bilateral price discovery protocol *BPD* can be specified to cover all possibilities. Listing 2 specifies a rule governing the `priceResponse` message.

The parameters are declared variables of their respective sorts. For each sort, relevant objects are declared, e.g., `DONE` and `FAILED` to indicate the result (`res`, `res1`, `res2`) in `priceResponse` and also in `executionConfirmation`. The variables `resID`, `reqID`, and others denote unique IDs for price response and request, respectively. Also, `execConfReq` can be `YES` or `NO` indicating whether or not a confirmation is required. The variable `ttl` indicates the valid time for the quoted rate as given in `rate`. As the request can be for two-way trades, the rate would typically include both a bid and an offer rate. We show one rate for simplicity. The `dir` in `priceAcceptance` indicates the direction: whether the *Taker* is buying or selling currency `cur1`. The amounts involved in the currency pair `cur1` and `cur2` are `amt1` and `amt2`, respectively. A disjunctive clause  $[\vee a \mid f(a)]$  with variable  $a$  ground to distinct objects  $a_i$  is equivalent to  $\bigvee_i f(a_i)$ .

### Listing 2. Specifying a rule for price response

```

1 nonexecutable act(priceResponse(m, t, resID, reqID,
2   res, execConfReq, ttl, rate))
3 if
4   -[ \ / cur1 \ / cur2 \ / amt1 |
5     fl(priceRequest(t, m, reqID, cur1, cur2, amt1))]
6 ++
7   (
8     ( fl(nothingDone(t, m, resIDa)) ++
9       fl(priceAcceptance(t, m, resIDb, dir))
10    )
11    &
12    -( fl(executionConfirmation(m, t, resIDb, FAILED)) &
13      fl(executionConfirmationAck(t, m, resIDb))
14    )
15  )
16 ++
17  act(priceResponse(m, t, resID, reqID, res2,
18    execConfReq2, ttl2, rate2))
19 ++
20  fl(priceResponse(m, t, resID, reqID, res,
21    execConfReq, ttl, rate))
22  where
23  res <> res2 ++ execConfReq <> execConfReq2
24  ++ ttl <> ttl2 ++ rate <> rate2.

```

The rule restricts occurrences of `priceResponse` if (a) no `priceRequest` with a matching `reqID` has happened, or (b) either a `nothingDone` has happened or a response has been already accepted, and the confirmation on that acceptance has not yet failed and any failure has not yet been acknowledged, or (c) a `priceResponse` with the same ID but a different result, confirmation requirement, `ttl`, or `rate` is happening simultaneously, or (d) a `priceResponse` with identical parameters has happened before. Here, (b) refers to confirmation as a `priceResponse` cannot happen again after it has been accepted, but can happen again if `executionConfirmation` for the accepted price fails and an acknowledgment for

this failure is sent (7.2.2, fourth, fifth, and sixth messages). Notice how lines 20–21, for instance, force us to answer the question of whether two priceResponses can have different confirmation requirements and ttl. The present formalization restricts a priceResponse only if one with the exact same parameters has already happened.

Listing 3 shows a specification of the nested commitment created as a result of a priceResponse. For brevity, only the case of confirmation not required is covered. A priceResponse creates the nested commitment (lines 1–3) as in cases (c), (d), (e), and (f) described in Section 3.1. However, to allow arbitrary levels of nesting, we substitute the inner commitments with placeholder conditions that are caused when the inner commitments are created. For example, lines 5–8 cause the precondition of the nested condition if the *Taker* has accepted to buying cur1 within ttl and commits to paying for it in the other currency. Similarly, lines 10–12 cause the condition of the nested commitment if the *Taker* has accepted to buy and the *Maker* commits to paying in the currency being bought. Similar rules would handle the case when the *Taker* is selling cur1.

**Listing 3. Specifying a nested commitment**

```

1 caused create(cc(m,t, priceResponsePrecond(resID),
2   priceResponseCond(resID))) if
3 act(priceResponse(m,t,resID,reqID,DONE,NO,ttl,rate)).

5 caused cond(priceResponsePrecond(resID)) if
6 fl(priceAcceptance(t,m,resID,TakerBuys)) &
7 -fl(ttlExpired(resID)) &
8 create(cc(t,m,T,pay(resID,amt2))).

10 caused cond(priceResponseCond(resID)) if
11 fl(priceAcceptance(t,m,resID,TakerBuys)) &
12 create(cc(m,t,T,pay(resID,amt1))).

```

Messages such as priceAcceptance and priceAcceptanceAck would cause the creation of the inner commitments depending on the interpretation adopted from Figure 3. Here, we interpret the meanings as in case (f). Also, a higher level of nesting can be modeled by having commitments as the conditions of the inner commitments. Rules for other messages and commitments can be specified similarly. Complete specifications are posted [1].

### 3.4 Querying the Specifications

The ability to query the formal specifications is crucial for discovering gaps, errors, and ambiguities. Protocol specifications can be compiled and queried via the causal calculator tool CCALC [18]. In essence, CCALC tries to find a model (a path in the transition system) that satisfies the constraints of the specification, given a query. The following describes several important queries that can help uncover problems in the price discovery specifications.

Listing 4 specifies a query to see if it can ever happen that one of the parties has a base commitment to another party

but there is no counter commitment that either currently exists or has been fulfilled. Note that unfulfilled conditional commitments are safe, but the same does not hold for base commitments. Thus, such a query should have no model in any protocol related to exchanges of any kind.

In Listing 4,  $p_i$  and  $q_i$  are variables of sort Condition. The label identifies this query and maxstep specifies the length of the history to be considered for search. Line 4 premises the query on the fact that initial holds in the starting state. The solver is asked to find the models for the formula of lines 5–10. A failure to find a model for this query is a necessary but not a sufficient condition to ensure commitment safety: a counter commitment may exist (resulting in a failure to find a model), but it may not be a commitment with the right condition.

**Listing 4. Querying for commitment safety**

```

1 :- query
2 label :: 1;
3 maxstep :: 0..infinity;
4 0: initial;
5 maxstep: (comm(cc(t,m,T,p1)) &
6   -[\q1 | comm(cc(m,t,T,q1)) ++ fl_cond(q1)]
7   ) ++
8   (comm(cc(m,t,T,p2)) &
9   -[\q2 | comm(cc(t,m,T,q2)) ++ fl_cond(q2)]
10  ).

```

Listing 5 specifies a query to see if the protocol does what is intended: get the deal done and end in a good state. The query formula represents the state at the end of case (f) with the case of *Taker* buying cur1. As fluents are inertial, such state queries are easy to formulate—it does not matter *when* the priceAcceptance happened, as long as it has happened in the history and the *Taker* has indicated to buy.

**Listing 5. Querying for successful execution**

```

1 :- query
2 label :: 2; maxstep :: 0..infinity; 0: initial;
3 maxstep: comm(cc(t,m,T,priceAcceptanceCond(resID))) &
4   comm(cc(m,t,T,pay(resID,amt1))) &
5   fl(priceAcceptance(t,m,resID,TakerBuys)).

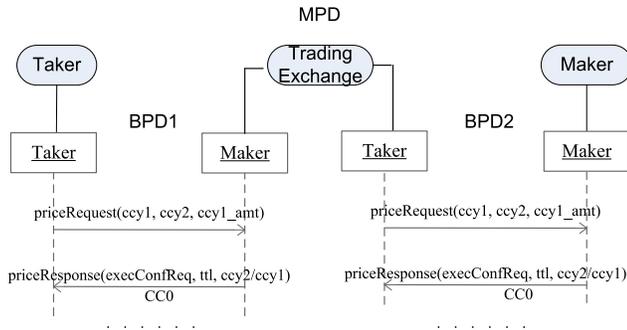
```

## 4 Composition of Protocols

Given a repository of modular protocol specifications for the core FX interaction patterns, it is natural to compose these modules to derive varieties of composite foreign exchange protocols as needed. The ability to reuse and compose existing protocols not only simplifies and improves engineering, but also provides new insights about the business processes. For example, the messages in TWIST process 7.2.3 are not new; they have already been described in processes 7.2.1 and 7.2.2. The technique of protocol composition introduced by Desai *et al.* [7, 8] enables defining process 7.2.3 in terms of processes 7.2.1 and 7.2.2.

Say, a bilateral price discovery protocol (*BPD*) is available. Figure 4 shows how multilateral price discovery (*MPD*) is specified by unioning two copies of *BPD* and stating constraints (known as *composition axioms*):

- $A_{X1}. MPD.Taker = BPD1.Taker$
- $A_{X2}. MPD.Maker = BPD2.Maker$
- $A_{X3}. MPD.TradeEx = BPD1.Maker, BPD2.Taker$
- $A_{X4}. BPD1.priceRequest.cur_1 \rightsquigarrow BPD2.priceRequest.cur_1$
- $A_{X5}. BPD1.priceRequest.cur_2 \rightsquigarrow BPD2.priceRequest.cur_2$
- $A_{X6}. BPD1.priceRequest.amt_1 \rightsquigarrow BPD2.priceRequest.amt_1$
- $A_{X7}. BPD2.priceResponse.execConfReq \rightsquigarrow BPD1.priceResponse.execConfReq$
- $A_{X8}. BPD2.priceResponse.rate \rightsquigarrow BPD1.priceResponse.rate$



**Figure 4. MPD by composing BPD with itself**

*Role definition* axioms (first three in the above) define a new role on the left in the composite protocol in terms of the roles of the component protocols on the right. As a result, the roles of the component protocols are renamed to be the new role. In this example, the trading exchange role mediates between the traditional taker and multiple makers.

*Data flow* axioms (next five in the above) specify that the parameter on the right gets its value from the parameter on the left. Thus, the message on the right cannot happen until all the parameters it needs have been bound (i.e., the suitable messages have happened). In this example, data flow axioms specify the constraint that the currency pair and the amount for which the trading exchange requests the maker must be identical to those received from the taker. Also, the confirmation requirement and the rate indicated by the maker to the trading exchange should be propagated to the taker. Thus, the trading exchange reduces to a simple mediator. A new rule per axiom is added to the theory to effect the binding of the parameters and the temporal ordering of the messages. The result of the composition would be the formal specification *MPD* of process 7.2.3, as posted [1]. Note that the resultant protocol *MPD* can be added to the protocol repository, and thus reused just like the core protocols. For example, Table 1 uses *MPD* (which is  $BPD \oplus$

Specification pattern	Protocols for pattern
7.2.1	<i>BPD</i>
7.2.2	<i>BPD</i>
7.2.3	$BPD \oplus BPD$
7.2.4 (order)	<i>Order</i>
7.2.5 (order, cancel)	<i>Order</i>
7.2.6 (credit check)	$BPD \oplus Credit$
7.2.7 (credit check–multi)	$BPD \oplus BPD \oplus Credit$
7.2.8 (price stream)	<i>ESP</i>
7.2.9 (price stream–multi)	$ESP \oplus ESP$

**Table 1. Mapping TWIST Sec. 7.2 processes**

*BPD*) with *Credit* to derive process 7.2.7. Additional kinds of composition axioms [7] are not needed here.

Various interesting business scenarios are possible depending on the composition axioms specified. Consider for example that  $A_{X8}$  were not specified. This would mean that the trading exchange could act as a secondary price maker and manipulate the bid-offer spread received from the primary maker. If  $A_{X7}$  were not specified, it would mean that the trading exchange could take risks of its own, and not require confirmation from the taker independent of the confirmation requirement indicated by the primary maker. Further, if  $A_{X6}$  were not specified, the trading exchange could either fill the requested amount from multiple makers or fill multiple taker requests from a single maker deal. Thus, composition axioms act as elegant, vivid specifications of configuration parameters. Modifying the axioms enables us to model vastly different business requirements. These possibilities are lost when informal specifications are constructed. The current text-based TWIST specification is ambiguous about these possibilities. Highlighting such possibilities is an important contribution of this paper.

Chapter 7 of the TWIST specification describes 28 interactions patterns. The above methodology helps model such patterns in terms of 12 core formally specified protocols and their compositions. More importantly, as demonstrated above, combinations *beyond* those described in the specification can be derived via novel compositions of the core protocols. Table 1 lists some of the patterns from TWIST Sec. 7.2, and shows how they can be modeled in terms of protocols. Here  $\oplus$  denotes the composition of the operand protocols. Four protocols *BPD*, *Order*, *Credit*, and *ESP* are enough to model nine patterns. Table 1 also points to interaction possibilities not covered by the TWIST specification. For example, each of the processes 7.2.4, 7.2.5, 7.2.8, and 7.2.9 may be composed with credit checks, if needed.

## 5 Discussion

The idea of modeling business processes based on conversation protocols is gaining interest. WS-CDL, a language for specifying such conversations among web services is being standardized by W3C [13]. Fu *et al.* specify conversation protocols as guarded automata [10]; studies of formal verification of conversations include [10, 14]. Zaha *et al.* propose focusing on the global view of interaction among services in an SOA to see if all the constraints of the global interaction can be enforced locally [23]. However, as Section 3.1 demonstrates, a contractual semantics is essential to characterize business interactions unambiguously. Other approaches to a contractual semantics, e.g., that of Davulcu *et al.* [5], lack the flexibility of commitments.

Singh *et al.* outline a vision for commitment-oriented modeling for engineering large-scale business processes [17]. Winikoff provides a set of guidelines for designing and implementing interactions based on commitments [21]. Desai *et al.* offer intuitions behind composition of commitment protocols [7]. The above works, however, are not sufficiently formalized to support verification as an integral engineering activity. The present paper builds on our recent formal approach for commitment protocols and their composition in  $C+$  [4, 6, 8]. The  $C+$  representation can be verified via the CCALC tool.

The above approaches, however, lack validation with respect to a practical case study. This is a key distinguishing feature of the present effort.

**Conclusions.** The broadest contribution of this paper is a methodology for the specification and engineering of business processes, applicable to domains where standardized business interactions are desirable. This methodology builds on a commitment-based representation of protocols, which captures business relationships among autonomous participants. The exercise of developing this representation identifies gaps and ambiguities in designers' understanding. Further, the resulting formal refactoring of the specifications is not only more compact, but also enables supporting a rich variety of business scenarios via composition, thus expanding the expressiveness of the representations.

At first sight, formal specifications might appear to demand more effort than sequence diagrams. However, formal specifications substitute not only for the diagrams but also for the informal descriptions that accompany them. Rigor and precision are indispensable, especially when diverse implementations have to interoperate in critical business processes. Further, high-level specifications can be validated with respect to requirements, thus yielding pay-offs in correctness and increased confidence.

Future work includes the development of graphical tools to simplify the specification and verification of protocols and their compositions.

## References

- [1] Price discovery protocol specifications in  $C+$ , 2007. <http://research.csc.ncsu.edu/mas/causal/>.
- [2] Aite Group. Electronic FX: welcome to the banks' neverland. <http://www.aitegroup.com/reports/200704201.php>.
- [3] Bank for International Settlements. Triennial central bank survey of foreign exchange and derivatives market activity. <http://www.bis.org/triennial.htm>.
- [4] A. K. Chopra and M. P. Singh. Contextualizing commitment protocols. In *AAMAS*, pages 1345–1352, 2006.
- [5] H. Davulcu, M. Kifer, and I. V. Ramakrishnan. CTR-S: a logic for specifying contracts in semantic web services. In *WWW Alt. Track Papers & Posters*, pages 144–153, 2004.
- [6] N. Desai, A. K. Chopra, and M. P. Singh. Representing and reasoning about commitments in business processes. In *AAAI*, 2007.
- [7] N. Desai, A. U. Mallya, A. K. Chopra, and M. P. Singh. Interaction protocols as design abstractions for business processes. *IEEE Trans. Soft. Engg.*, 31(12):1015–1027, 2005.
- [8] N. Desai and M. P. Singh. A modular action description language for protocol composition. In *AAAI*, 2007.
- [9] FIX Protocol Ltd. (FPL). FIX 5.0 protocol specifications. <http://tinyurl.com/auzx9>.
- [10] X. Fu, T. Bultan, and J. Su. Conversation protocols: A formalism for specification and verification of reactive electronic services. *Theoret. Comp. Sci.*, 328(1–2):19–37, 2004.
- [11] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153(1–2):49–104, 2004.
- [12] International Swaps and Derivatives Association (ISDA). FpML 4.1 recommendation. <http://tinyurl.com/3yle3u>.
- [13] N. Kavantzaz, D. Burdett, G. Ritzinger, T. Fletcher, and Y. Lafon. Web services choreography description language 1.0 (w3c recommendation), 2005. <http://www.w3.org/TR/ws-cdl-10>.
- [14] R. Kazhamiakin and M. Pistore. Static verification of control and data in web service compositions. In *ICWS*, pages 83–90, 2006.
- [15] RosettaNet. Home page. [www.rosettanet.org](http://www.rosettanet.org), 1998.
- [16] M. P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7:97–113, 1999.
- [17] M. P. Singh, A. K. Chopra, N. Desai, and A. U. Mallya. Protocols for processes: programming in the large for open systems. *ACM SIGPLAN Notices*, 39(12):73–83, 2004.
- [18] Texas Action Group at Austin. The causal calculator CCALC. <http://www.cs.utexas.edu/users/tag/cc/>.
- [19] Transaction Workflow Innovation Standards Team (TWIST). TWIST wholesale trade lifecycle. <http://tinyurl.com/yswnqq>.
- [20] Wikipedia. Foreign exchange markets. [http://en.wikipedia.org/wiki/Foreign\\_exchange\\_market](http://en.wikipedia.org/wiki/Foreign_exchange_market).
- [21] M. Winikoff. Implementing commitment-based interaction. In *AAMAS*, 2007.
- [22] P. Yolum and M. P. Singh. Enacting protocols by commitment concession. In *AAMAS*, 2007.
- [23] J. M. Zaha, M. Dumas, A. ter Hofstede, A. Barros, and G. Decker. Service interaction modeling: Bridging global and local views. In *EDOC*, pages 45–55, 2006.