

Interim Report

IR-98-115

Modeling a Decentralized Asset Market: An Introduction to the Financial “Toy-Room”

Francesca Chiaromonte (chiaro@iiasa.ac.at)

Giovanni Dosi (dosi@iiasa.ac.at)

Approved by

Gordon J. MacDonald (macdon@iiasa.ac.at)

Director, IIASA

October 1998

Acknowledgements

We would like to thank M. Bertè, who implemented the code for Toy-Room, and helped us with the definition of several parts of the model. We would also like to thank A. Bassanini, H. Tordjman, and all members of the TED project at I.I.A.S.A. who offered comments and insights at various stages of the development of this work.

Abstract

In this paper, we describe a micro-founded simulation environment for decentralized trade in a financial asset. Within the philosophy of computer-simulated “artificial markets”, this environment allows one to experiment in a modular fashion with (i) individual characterizations in terms of behaviors and learning, (ii) different architectural and institutional traits of the market, and (iii) time-embedding of events at the system and the individual level.

Contents

1	General description	2
1.1	The board	4
1.2	The Phones	5
1.3	Transactions	6
1.4	The Traders	7
2	The Dynamics	10
2.1	Concluding transactions: the Trading Round	10
2.2	Completing transactions	12
2.3	Updating the trading documents: Targeting/Positioning	13
2.4	Expulsion: leaving the room	13
2.5	Entering the room	15
2.6	The “external value(s)”	17
3	Behavioral repertoires and market features	17
3.1	The honoring constraint	18
3.2	Synchronous vs asynchronous completion markets	18
3.3	Imitation	19
3.4	(Functional) differentiation of traders	20
3.5	A tentative behavioral taxonomy for spot-spot trading based on “price assessment”	20
3.6	Evolution of behaviors	25
4	The Statistical Office	25
5	Conclusions	27

Modeling a Decentralized Asset Market: An Introduction to the Financial “Toy-Room”

Francesca Chiaromonte (chiaro@iiasa.ac.at)

Giovanni Dosi (dosi@iiasa.ac.at)

Introduction

In this paper, we describe a micro-founded simulation environment –the Financial “Toy-Room” (FTR)– for decentralized trade in a financial asset. Some aspects of the representation are intentionally kept very simple, and in a sense abstract: quite diverse models may indeed be implemented as particular instantiations of the general template presented in the following.

The general motivations for FTR are to a good extent akin to those inspiring already existing computer-simulated “artificial markets” of a financial asset, such as those by Marengo and Tordjman (1996), Rieck (1994), Beltratti and Margarita (1992), and Arthur et al. (1997).

Obvious common points of departure are (i) the acknowledgment of the limitations of models of market dynamics centered upon the behavior of a mythical representative agent endowed with unbiased forward-looking expectations, and conversely (ii) the challenge of nesting the theory into an explicit account of heterogeneous, interacting agents.

Some forms of heterogeneity in information and beliefs can be incorporated into analytically tractable models (see for example the information-related heterogeneity in Grossman and Stiglitz, 1976 and 1980, the diversity of beliefs associated to the presence of “noise” traders in De Long et al. 1990, 1991 and Schleifer and Summers, 1990, see also Blume and Easley, 1990). However, analytical tractability poses heavy constraints on the forms and degrees of heterogeneity, as well as the forms of learning, one can handle. Moreover, one is forced to analyze almost exclusively limit (equilibrium) properties of the models, and neglect finite-time properties which might nonetheless be the most relevant for comparison with empirical data.

The “artificial market” approach tries to overcome these drawbacks by explicitly simulating populations of interacting agents who might endogenously evolve beliefs, behaviors and “mental models” (Marengo and Tordjman, 1996): FTR has been build on the grounds of the same basic philosophy. At the same time FTR, when compared to other “artificial markets”, enlarges the scope of analysis in several respects.

First, FTR entails easy experimentation with different types of agents, both in terms of behavioral and cognitive patterns, and in terms of learning procedures.

Second, it allows exploration of the properties of different architectural and institutional traits, especially with respect to the “physics” of interactions (e.g. the specific mechanism for decentralized encounters), and the information availability by individual traders –or groups of them.

Third, FTR embodies an explicit time-embedding of events that allows us to easily represent asynchronous and/or diversely paced “clocks” for diverse classes of events at the system and individual level (e.g. buying and selling –trading, vs. accessing “news”, vs. making trading decisions, vs. learning). Relatedly, FTR naturally allows us to study the dynamic properties of the system on different time-scales.

As such, we see FTR as the “artificial” counterpart of micro-structural studies (cf. Frankel, Galli and Giovannini, 1997 and Goodhart and Payne, 1996). There is a long and growing list of “stylized facts” to a good extent still in search of an interpretation (for complementary discussions, see Brock, 1997, Frankel, Galli and Giovannini, 1997, Goodhart and Figliuoli, 1991, Guillaume et al. 1997). With FTR, one can investigate what types of cognitive/behavioral patterns and learning processes, and what types of interaction and information regimes, can reproduce the regularities detected in empirical markets as emergent properties of the artificial market dynamics.

A second class of exercises, although partly overlapping with the above, have the primary nature of *thought experiments* on the effect of individual characterization and institutional set-up upon system dynamics. Two broad questions come immediately to mind, namely:

1. Holding individual characteristics (i.e. cognitive/behavioral patterns, and possibly learning processes) and information regime constant, what happens if one changes the interaction regime?
2. Holding the institutional set-up (i.e. interaction and information regimes) constant, what happens as one varies the “ecology” of cognitive/behavioral patterns and learning processes?

In connection with empirical studies, simulation experiments will allow us to assess whether observed statistical regularities (e.g. the so called “ARCH” effects, “fat tails”, etc.) are generic properties, holding over a wide range of interaction regimes and “ecologies”, or conversely, whether such regularities are conditional to very specific institutional set-ups and distributions of agents’ “types”.

Details on the computer implementation of FTR are given in Bertè (1998), and some preliminary simulation experiments are reported in Chiaromonte and Bertè (1998).

All through the paper, we stress modularity and comment extensively on how various components of the environment can be used, modified or extended while maintaining the general framework. Section 1 describes the structure of FTR, and the main entities in it. Section 2 describes the dynamics; that is, how the entities and the variables associated with them may evolve over time. Section 3 provides some illustrative examples of how individual behaviors may be specified and, together, of alternative trading scenarios. Section 4 concerns the collection of simulation outputs. Conclusions are given in Section 5.

1 General description

Let us begin with a somewhat loose but intuitive introduction to our artificial market and its basic building-blocks. Our metaphor for the market is a *room*, inhabited by actual and would-be traders, and provided with both displays of information on what goes on in the market itself,

and communication lines with the outside world (which represent mechanism of observation of purportedly “fundamental” and non-fundamental underlying economic variables).

In this metaphor, *where* in the room traders position themselves maps into micro-decisions –e.g. *seeking* or *accepting* transactions under certain price ranges, remaining *inactive*, etc.

As already hinted above, the basic philosophy of FTR entails a modular separation of (i) the “physics” of interactions among traders and the rules by which trade takes place, (ii) the information traders might access, (iii) the algorithms by which traders process such information in making decisions, and eventually the algorithms by which they learn (i.e. evolve their beliefs and decision-making procedures).

In our metaphor, traders are fully described by their “trading documents”, their “note-pads”, and their “manuals”.

Trading documents are a sort of “identity card” of the trader at any particular time, reporting his disposition to, for example, seek/accept a selling/buying transaction (captured by 0-1 *flags*), and the prices or price-spreads at which the transaction is sought/acceptable. Moreover, since traders show each other their documents (or parts of them) upon meeting, these “identity cards” vector information in pair-wise encounters.

The *note-pad* contains the “internal-memory” of the trader, recording, for example, the sequence of transactions he undertook in the past, and information on other traders.

The *manual* embodies decision and learning algorithms, which of course might range from simple technical rules to sophisticated calculating abilities.

The *board*, on one wall of the room, displays all information on market dynamics publicly available to traders. Moreover, the board displays signals by which traders are called upon participating in pair-wise encounters, and signals concerning the time-scansion.

Finally, *phones* stand for access to outside information (i.e. to “news” concerning fundamental, and possibly non-fundamental, variables). Access can be unlimited, or restricted to a subset of traders, as well “toll-free” or costly.

Note that the information regime is defined by what is reported on trading documents (or more specifically, what parts of the documents traders are required to show each other upon meeting), what is placed on the board, and what goes through phone lines (specifying accessibility, and possibly fees).

Given this overview, let us move to a more detailed description of FTR. For the time being, we assume the asset to be homogeneous. The room is inhabited by a group of *traders* $T \in \mathcal{T}$ engaging in *transactions* $o \in \mathcal{O}$. Along the first wall, there is a row of *windows*. Along the second wall, there is a row of *chairs*. Along the third wall, there is room for by-standers and a *door* through which traders enter and leave the room.

A trader’s position is expressed via the values of some 0-1 flags (see trading documents below)

$$f^\alpha[T] = \max\{f_b^\alpha[T], f_s^\alpha[T]\} \quad , \quad f^\sigma[T] = \max\{f_b^\sigma[T], f_s^\sigma[T]\} \in \{0, 1\}$$

T can stand by the third wall ($f^\alpha[T] = f^\sigma[T] = 0$), be behind a window ($f^\alpha[T] = 1, f^\sigma[T] = 0$), in a chair ($f^\sigma[T] = 1, f^\alpha[T] = 0$), or behind a window and in a chair simultaneously ($f^\alpha[T] = f^\sigma[T] = 1$). Standing by the third wall, a trader renounces involvement. Behind a window, a

trader is in the role of *acceptor* of transactions. In a chair, he is in the role of *seeker*. These embody two different attitudes towards trading that we wish to superimpose to the buying/selling distinction. As we will see, seekers are designed to be the active parties, and only some particular traders might be allowed to hold a window and a chair at the same time.

1.1 The board

On the fourth wall, there is a board through which one governs encounters among traders, time representation, and flows of public information. The board contains a number of “slots”, namely:

- Two *callers*, one for acceptors and one for seekers

$$\begin{aligned}\alpha &\in \{T \in \mathcal{T} : f^\alpha[T] = 1\}^k \quad (k \in \mathbb{N}^1) \\ \sigma &\in \{T \in \mathcal{T} : f^\sigma[T] = 1\}\end{aligned}$$

which are used to implement encounters among traders (see below).

- A *transaction counter* $N = \text{card}(\mathcal{O}) \in \mathbb{N}^{1-1}$.
- A *clock* ticking minutes ²

$$H \in \mathbb{N}^1 \text{ for } \sum_{j=1}^{H-1} \nu_j \leq N < \sum_{j=1}^H \nu_j$$

We name $\nu_1, \nu_2, \dots \in \mathbb{N}^1$ *system converting sequence*. It converts “time in terms of transactions” into “time in system-minutes”. The issue here is how to translate a time pace defined in terms of transactions into some sort of “objective” time for system-level events –which are linked to the board clock, as well as some sort of “internal” time for trader-specific events –which are linked to individual watches (see below). The system converting sequence allows us, among other things, to represent accelerations and decelerations of the trading process in system-minutes. In the following, we will often refer to the number of transactions in a minute as its *length-in-transactions* ³.

- A *display*, reporting (public) information of various kinds:
 - A *tape* showing all transaction prices $p[o] \in \mathbb{R}_+^1$ up to the latest, in the order in which they occurred, say

$$p(1), p(2), \dots, p(N)$$

- A *disclosure sheet* containing the names and current asset levels of all traders whose asset endowment exceeds a given threshold

$$(T, q[T]) \in \mathcal{T} \times \mathbb{N}^1, \quad \forall T : q[T] \geq Q \in \mathbb{N}^1$$

¹The reference is to concluded transaction, regardless of whether they have completed yet (see below). $\text{card}(\cdot)$ indicates the cardinality; that is, the number of elements of the argument set.

²A sum on an empty set is assumed to be equal to 0.

³In fact, if the ν_j 's are different, equal durations in minutes can correspond to different lengths-in-transactions

This can be interpreted as an approximate representation of the requirement, typical of some stock markets, to disclose ownership of an asset (and/or bids for it) when exceeding a certain level (share). Thus, the threshold Q is an architectural parameter of the market capturing the extent of publicly available information about the “relative control” on the asset.

Notice that the board display is the locus for representation and management of *public* information flows: Any other publicly available information one might wish to introduce should be placed here.

1.2 The Phones

Again on the fourth wall, there are phones used by traders to obtain information from outside the room. In particular, phones convey information on variables affecting (or more generally related to) the “real” value of the object which the financial asset denominates. Such variables are assumed to be *independent of the trading process* (e.g. dividends on a stock, or other outside “news”). In the following, we will consider a single variable and simply refer to it as the *external value*.

- Through a first phone number, a trader T can access (a possibly *noisy version of*) the *current external value*

$$Z_{(H)} + \epsilon[T]$$

The current external value $Z_{(H)} \in \mathbb{R}_+^1$ evolves exogenously on the minute-scale (following the board clock). The noise $\epsilon[T]$ is a draw from a probability distribution (usually, but not necessarily, a 0-mean normal), and might be trader-specific ⁴. This could mean that different traders observe independent draws from the same distribution \mathcal{E} , or that they observe independent draws from different distributions $\mathcal{E}[T]$. We also keep the option that all traders dialing the first phone number within minute H on the board clock, observe a *single draw*, say $\epsilon_{(H)}$, from \mathcal{E} ⁵. Obviously, the noise can be eliminated by setting the distribution(s) variance(s) (and possibly mean(s)) to 0 ⁶.

- Through a second phone number, a trader can access *past external values*, say ⁷

$$Z_{(H-1)}, Z_{(H-2)} \dots$$

With this set-up, one might experiment with imperfect and asymmetric information. So, for example, one can assume that a subgroup of traders has access the history of the external value, while another subgroup has access to the (noisy) current value. In other words, phone numbers might not be known to all traders. Moreover, the numbers might be taken to be toll-free, or *fees* might be associated to them (i.e. information might be costly, as in Grossman and Stiglitz, 1980).

Phones are the locus for representation and management of *external* information flows: Any other information regarding variables independent of the trading process one might wish to

⁴Note that this permits the implementation of “noise” traders in the sense of De Long et al. (1991).

⁵(Common) draws will still be independent across H 's.

⁶To better understand the use of $\epsilon[T]$, we refer the reader to Chiaromonte and Bertè (1998), and Bertè (1998).

⁷In practice, the time series will be truncated a certain number of minutes “backwards”.

introduce, should go through phone lines for which the experimenter must specify accessibility, and possibly fees.

The extension to the case of many external values is straightforward; in our metaphor, it is just a matter of multiplying phone numbers. Traders could observe them alternatively or jointly. Moreover, one could distinguish between *fundamental variables*, that are indeed related to the “real” value of the object denominated by the asset, and *sun-spots*, that are not related to the “real” value, but are still observed and used by some traders in their decision-making (more details are given in Section 3.5).

1.3 Transactions

For the time being and for the sake of simplicity, we assume each transaction to concern only *one unit* of the asset. When formalizing decision-making by traders, this allows us to neglect quantities, and concentrate on prices and completion schedulings (see below). However, this represents a strong constraint that we plan to remove in the near future. In fact, limiting each transaction to one unit of asset, besides eliminating a crucial dimension of decision-making, has other implications due to its “interaction” with other features of FTR: fixing the system converting sequence, i.e. the number of transactions per minute, one fixes also the trading volume per minute

Transactions might or might not be spot. The *conclusion* of a transaction, i.e. the agreements on payment and delivery between two traders, might or might not coincide in time with its *completion*, i.e. the actual exchange of cash and asset unit.

To handle the time profile we associate to each transaction, together with the minute on the board clock during which it was concluded, $h[o] \in \mathbb{N}^1$, a *completion scheduling*

$$(dh_1[o], dh_2[o]) \in \mathbb{N}^2$$

and *completion flags*

$$(c_1[o], c_2[o]) \in \{0, 1\}^2$$

$h[o] + dh_1[o]$ and $h[o] + dh_2[o]$ express, respectively, the minutes on the board clock for the payment (from the buyer to the seller) and the delivery (from the seller to the buyer). $c_1[o]$ or $c_2[o]$ equal to 1 express, respectively, the fact that the payment or the delivery have occurred. We will often refer to transactions which have $c_1[o]$ and/or $c_2[o]$ equal to 0 as *outstanding*.

$dh_1[o] = 0$ is meant to represent a spot payment: the buyer $b[o]$ pays $p[o]$ to the seller $s[o]$ simultaneously to the transaction conclusion, whenever this occurred within $h[o]$. $c_1[0]$ will be 1 from the very start. On the other hand, a given $dh_1[o] > 0$ bounds $b[o]$ to pay $p[o]$ any time during minute $h[o] + dh_1[o]$ on the board clock⁸. At the beginning of $h[o] + dh_1[o]$ ⁹ the transaction enters its *completion phase for the buyer*, which will terminate when the payment occurs. $c_1[o]$, which was initialized at 0 upon conclusion, will then be set to 1. The completion phase is supposed to last at most one minute, regardless of the length-in-transactions of $h[o] + dh_1[o]$. However, we will see that it can be prolonged, even though not indefinitely and with a penalty for the procrastinating trader (see below on bonus-minutes).

⁸As we will see, completion is organized in such a way that both payments and deliveries due at certain minute are performed as soon as possible within that minute; that is, as soon as the involved traders have the necessary cash or asset units.

⁹This coincides with the board transaction counter showing $\sum_{j=1}^{h[o]+dh_1[o]-1} \nu_j$.

Similarly, $dh_2[o] = 0$ represents a spot delivery: the seller $s[o]$ delivers one unit of asset to the buyer $b[o]$ simultaneously to the transaction conclusion, whenever this occurred within $h[o]$. $c_2[0]$ will be 1 from the very start. $dh_2[o] > 0$ bounds $s[o]$ to deliver any time during minute $h[o] + dh_2[o]$. At the beginning of $h[o] + dh_2[o]$ the transaction enters its *completion phase for the seller*, which will terminate when the delivery occurs. $c_2[o]$, which was initialized at 0 upon conclusion, will then be set to 1 (again, the completion phase ought to last at most one minute but can be prolonged, even though not indefinitely and with a penalty).

It is important to remark that the two scheduling terms $dh_1[o]$ and $dh_2[o]$ need not coincide. Also, the lengths-in-transactions of the two completion phases can differ. While the scheduling is under traders' control (see below), different lengths of the completion phases could be due, besides traders' procrastination, to differences in the ν_j 's of the (system) converting sequence –that is, to the fact that some system minutes contain more transactions than others.

Transactions can be classified according to their completion scheduling as: (i) *spot-spot* ($dh_1[o] = 0$, $dh_2[o] = 0$), (ii) *short on the buying side* ($dh_1[o] > 0$, $dh_2[o] = 0$), (iii) *short on the selling side* ($dh_1[o] = 0$, $dh_2[o] > 0$), and (iv) *forward* ($dh_1[o] > 0$, $dh_2[o] > 0$). Thus, the model allows us to represent spot trading, short buying or selling, and forward trading.

1.4 The Traders

Each trader in the room is characterized by:

- A *flag for expulsion* (i.e. institutionally sanctioned bankruptcy) $ex[T] \in \{0, 1\}$. This flag is initialized at 0; if and when it is switched to 1, the trader is irreversibly removed from the room.
- A *counter of available bonus-minutes* $B[T] \in \mathbb{N}^I$. A certain number $B[T] = B_{max}$ of bonus-minutes is given to each trader when he enters the room. Those minutes are then used to extend completion phases of non spot-spot transactions in which the trader is involved (i.e. to postpone deliveries and/or payments with respect to the agreed upon scheduling), when needed. $B[T]$ decreases accordingly.

Note that this bonus-system constitutes an architectural trait of the market (B_{max} is an architectural parameter), which can be interpreted as a loose proxy for a *credit* system, which is not explicitly modeled in the current version of FTR ¹⁰ (more on the role of bonus-minutes will be given in Section 2.4).

- *Cash* and *asset* endowments; $m[T] \in \mathbb{R}_+^I$, $q[T] \in \mathbb{N}^I$.
- An indicator of what we shall call *behavioral state* $r[T] \in \{1, \dots, R[T]\}$ ($R[T] \in \mathbb{N}^I$), which captures the kind of algorithms used in decision-making (see description of the manual, below).

Moreover, the trader carries:

- *Trading documents*, in the form of an *acceptor* and a *seeker sheet*. Each sheet reports two flags, two reference prices, and two sets of completion scheduling options, respectively for

¹⁰The interpretation is straightforward when bonus-minutes are used to postpone payments, and less immediate, but similar, when they are used to postpone deliveries –as if traders could borrow from a “bank” asset units, as well as cash.

buying and selling:

$$f_b^\alpha[T], f_s^\alpha[T] \in \{0, 1\} \quad , \quad p_b^\alpha[T], p_s^\alpha[T] \in \mathbb{R}_+^I \quad , \quad D_b^\alpha[T], D_s^\alpha[T] \subseteq \mathbb{N}^2$$

$$f_b^\sigma[T], f_s^\sigma[T] \in \{0, 1\} \quad , \quad p_b^\sigma[T], p_s^\sigma[T] \in \mathbb{R}_+^I \quad , \quad D_b^\sigma[T], D_s^\sigma[T] \subseteq \mathbb{N}^2$$

If $f_b^\alpha[T] = 1$, T is accepting transactions as a buyer, at prices $p \leq p_b^\alpha[T]$ and completion schedulings $(dh_1, dh_2) \in D_b^\alpha[T]$. If $f_s^\alpha[T] = 1$, T is accepting transactions as a seller at prices $p \geq p_b^\alpha[T]$ and completion schedulings $(dh_1, dh_2) \in D_s^\alpha[T]$. Similarly, if $f_b^\sigma[T] = 1$, T is seeking transactions as a buyer, at prices $p \leq p_b^\sigma[T]$ and completion schedulings $(dh_1, dh_2) \in D_b^\sigma[T]$, while if $f_s^\sigma[T] = 1$, T is seeking transactions as a seller, at prices $p \geq p_b^\sigma[T]$ and completion schedulings $(dh_1, dh_2) \in D_s^\sigma[T]$.

Clearly, the variables in the trading documents constitute the *main decision variables* for the trader. In the following, we call *positioning* the determination of flags, and *targeting* the determination of the other components of the trading documents (reference prices and completion scheduling options).

As we will see, traders are required to show each other their acceptor or seeker sheets (or parts of them) when encountering. Thus, trading documents convey *pair-wise information exchanges*. One might want to introduce some form of censoring (e.g. traders, or some subgroup of them, might be assumed to disclose their willingness to buy or sell, but not their reference prices) in order to capture different institutional rules on information disclosure.

Trading documents, possibly with censoring, are the locus for representation and management of *pair-wise* information flows: Any other information that one might want to be exchanged by traders upon meeting each other should be placed here.

- A *watch* ticking minutes

$$H[T] \in \mathbb{N}^I \quad \text{for} \quad \sum_{j=1}^{H[T]-1} \nu_j[T] \leq N < \sum_{j=1}^{H[T]} \nu_j[T]$$

where $\nu_1[T], \nu_2[T], \dots \in \mathbb{N}^I$ is T 's *converting sequence*. It converts “time in terms of transactions” into “time in T -minutes”; that is, into some sort of “internal” time for trader-specific events –which are linked to individual watches. We use this in representing traders’ decision-making processes, and possibly modification of decision algorithms and learning.

The watch is not necessarily synchronized with other traders’ watches, or with the board clock, in the sense that the converting sequences might differ. Thus, T 's “internal” watch-time might be unrelated to that of other traders, and to “objective” board clock-time.

The nature and relations among system and traders’ converting sequences can be interpreted as both architectural traits of the the market, and behavioral characteristics of traders. Let us mention a few simple instances: the system sequence could be one of fixed numbers, all equal to each other (all minutes on the board clock have the same length-in-transactions). Alternatively, the system sequence could be a sequence of independent draws from a given distribution \mathcal{N} on \mathbb{N}^I . Traders’ sequences could just all copy the system one $\nu_j[T] = \nu_j, \forall j = 1, 2, \dots, \forall T \in \mathcal{T}$, or be otherwise fixed. Also, traders’ sequences could be themselves sequences of independent draws from distributions $\mathcal{N}[T]$ on \mathbb{N}^I , and these distributions could be taken to coincide with \mathcal{N} , or be given otherwise.

- A *note-pad* reporting (private) information of various kinds:

- A *record for each transaction* the trader has concluded, with the identities of buyer and seller (e.g. $b[o] = T$), the transaction price, the transaction time, completion scheduling, and completion flags

$$\begin{aligned} b[o] \in \mathcal{T} \quad , \quad s[o] \in \mathcal{T} \quad , \quad p[o] \in \mathbb{R}_+^1 \\ h[o] \in \mathbb{N}^1 \quad (dh_1[o], dh_2[o]) \in \mathbb{N}^2 \quad , \quad (c_1[o], c_2[o]) \in \{0, 1\}^2 \end{aligned}$$

- A *record for each other trader* $T' \neq T$ in the group he has encountered, with the acceptor and seeker sheets of T' as they appeared upon the last encounter, and the time of such encounter

$$\begin{aligned} \hat{f}_b^\alpha(T')[T], \hat{f}_s^\alpha(T')[T] \quad , \quad \hat{p}_b^\alpha(T')[T], \hat{p}_s^\alpha(T')[T] \\ \hat{D}_b^\alpha(T')[T], \hat{D}_s^\alpha(T')[T] \quad , \quad h^\alpha(T')[T] \in \mathbb{N}^1 \end{aligned}$$

$$\begin{aligned} \hat{f}_b^\sigma(T')[T], \hat{f}_s^\sigma(T')[T] \quad , \quad \hat{p}_b^\sigma(T')[T], \hat{p}_s^\sigma(T')[T] \\ \hat{D}_b^\sigma(T')[T], \hat{D}_s^\sigma(T')[T] \quad , \quad h^\sigma(T')[T] \in \mathbb{N}^1 \end{aligned}$$

Again, notice that parts of this information might be censored. Moreover, one could easily introduce a form of time-decay, i.e. progressively remove records relative to encounters that date more than a given number of minutes backwards.

- A *manual* containing algorithms which embody the trader’s behavioral repertoire. As already mentioned, the manual is, so to speak, the “brain” wherein rests all behavioral and *latu sensu* “cognitive” attributes one gives to the trader. As it stands now, the manual has two chapters:

- Chapter 1: *Targeting/positioning algorithms*. These algorithms are used to update the variables in the acceptor and seeker sheets, and thereby also the position of the trader in the room.
- Chapter 2: *Transaction-selection algorithms*. These algorithms are used when seeking transactions, to decide which to conclude among the ones made available by acceptors.

In turn, each chapter contains $R[T]$ alternative sets of algorithms to be used, respectively, when the behavioral state is $r[T] = 1$, $r[T] = 2$, etc. In other words, what we call a behavioral state can be seen as a collection of behavioral/cognitive patterns relative to the various tasks addressed by Chapters 1 and 2 of the manual. As we will see in detail in Section 3, behavioral states can be used in a variety of ways; just to mention some examples, one $r[T]$ might correspond to being a fundamentalist, while another $r[T]$ might correspond to being a particular type of chartist. Yet another $r[T]$ might correspond to specific behavioral patterns followed while trying to cover open positions (i.e. during completion phases of transactions the trader is involved in).

Within our metaphor, whenever traders can switch between behavioral states, switching rules could be placed in a third chapter of the manual.

Moreover, a fourth chapter of the manual will eventually preside over the evolution of the behavioral/cognitive patterns themselves (e.g. through processes of experimentation and inductive adaptation similar to those modeled by Marengo and Tordjman, 1996, or Arthur et al. 1997 –see also Section 3.6).

Clearly, Chapters 1 and 2 on one side, and Chapters 3 and 4 on the other, have a different role and nature: the former contain algorithms to trade, while the latter contain “higher

level” algorithms to switch between, or evolve, the previous ones. In the following, we use the word manual (space of manuals, etc.) to refer to Chapters 1 and 2.

2 The Dynamics

2.1 Concluding transactions: the Trading Round

Let us now describe a standard *trading round*, which might or might not produce an actual transaction. As we will see, the trading round specification embodies all rules concerning who trades with whom, and how.

The seeker caller on the board switches on and shows the name $\sigma = T$ of a trader drawn at random among the ones waiting in chairs (i.e. such that $f^\sigma[T] = 1$). T leaves his chair. The acceptor caller on the board switches on and shows the names $\alpha = \{T'_1, \dots, T'_k\}$ of k traders drawn at random among the ones behind windows (i.e. such that $f^\alpha[T'] = 1$), T itself excluded (in case he had both flags equal to 1, i.e. was in a chair and at a window simultaneously). These are the acceptors the seeker has access to. Clearly, seeker and acceptors involved in the round could be identified with procedures other than (uniform) random drawing.

T approaches all $T' \in \alpha$ at their windows. In each approach, acceptor and seeker are required to show each other their acceptor and seeker sheets. Hence, both update the other’s record in their note-pads, with $h^\alpha(T')[T] = h^\sigma(T)[T'] = H$, the current minute on the board clock. After having collected the information, T must decide what to do. Suppose $f_b^\sigma[T] = 1$. Then, a first set of transactions that are available to T is represented by:

$$\begin{aligned} b[o] = T, \quad s[o] = T' \quad , \quad p[o] = \lambda p_b^\sigma[T] + (1 - \lambda) p_s^\alpha[T'] \\ h[o] = H \quad , \quad (dh_1[o], dh_2[o]) \sim Un(D_b^\sigma[T] \cap D_s^\alpha[T']) \end{aligned}$$

for each T' (among the k acceptors) such that $f_s^\alpha[T'] = 1$, $p_b^\sigma[T] \geq p_s^\alpha[T']$ and $D_b^\sigma[T] \cap D_s^\alpha[T'] \neq \emptyset$. Furthermore, if $f_s^\sigma[T] = 1$, a second set of transactions that are available to T is represented by:

$$\begin{aligned} b[o] = T', \quad s[o] = T \quad , \quad p[o] = \lambda p_s^\sigma[T] + (1 - \lambda) p_b^\alpha[T'] \\ h[o] = H \quad , \quad (dh_1[o], dh_2[o]) \sim Un(D_s^\sigma[T] \cap D_b^\alpha[T']) \end{aligned}$$

for each T' (among the k acceptors) such that $f_b^\alpha[T'] = 1$, $p_s^\sigma[T] \leq p_b^\alpha[T']$ and $D_s^\sigma[T] \cap D_b^\alpha[T'] \neq \emptyset$. The symbol $Un(\cdot)$ indicates a uniform probability distribution on the elements of the argument set, and the draws generating $(dh_1[o], dh_2[o])$ for each of the available transactions are taken to be independent. Again, completion schedulings could be determined in ways other than independent (uniform) random drawing from the set of completion scheduling options that are common to the traders involved. A form of non-random determination, parameterized through $\lambda \in [0, 1]$, is given for the price. Also price determination could be implemented in a different fashion.

$\lambda \in [0, 1]$, and $(dh_1[o], dh_2[o])$ for each available transaction, are supposed to be known to T . Moreover, the initial completion flags for any of the available transactions would be set to:

$$\begin{aligned} c_1[o] = 1 \quad \text{if } dh_1[o] = 0, \quad c_1[o] = 0 \quad \text{otherwise} \\ c_2[o] = 1 \quad \text{if } dh_2[o] = 0, \quad c_2[o] = 0 \quad \text{otherwise} \end{aligned}$$

If the overall set of available transactions is empty, seeker and acceptor callers on the board switch off, and the round ends without the conclusion of a transaction.

Suppose now the overall set of available transactions is not empty. Then T selects one among them using the transaction-selection algorithms in his manual (which might differ depending on his behavioral state). He goes to the corresponding T' , and the two conclude the transaction. A new o is added to \mathcal{O} , the board transaction counter shifts by 1, and the transaction price $p[o]$ is appended as the latest price on the board display tape.

A record of o is added to T and T' note-pads. Moreover, if the transaction is spot on at least one side, the cash and/or asset levels are updated right away. For example, taking the case $b[o] = T$:

$$\begin{aligned} m[T] &\leftarrow m[T] - p[o] , & m[T'] &\leftarrow m[T] + p[o] & \text{if } dh_1[o] = 0 \\ q[T] &\leftarrow q[T] + 1 , & q[T'] &\leftarrow q[T] - 1 & \text{if } dh_2[o] = 0 \end{aligned}$$

In the case of spot delivery, the updated levels of asset of both traders are checked to determine whether T and/or T' must be added to, or removed from, the disclosure sheet in the board display.

The board clock, as well as the watches of all traders in the room, might or might not shift by 1 (depending on the system and traders' converting sequences). Seeker and acceptor callers on the board switch off. With both callers switched off, the room is ready to undergo the next trading round.

k is an architectural parameter of the market: it represents the size of the sample of acceptors that a seeker has access to in one round. If $k = 1$ the seeker scans a single acceptor. Hence, he collects "fresh" information on only one trader, and the transactions available to him might be none, one, or at most two (one involving him as a buyer, and one involving him as a seller). At the opposite extreme, if $k \geq \text{card}(T' \in \mathcal{T} : f^\alpha[T'] = 1)$ the seeker scans everyone who is willing to accept transactions (except possibly himself!). Hence, he collects "fresh" information on all acceptors, and can choose among all potential transaction partners. The interpretation in terms of degrees of "informational perfection" and "globality" of interactions is straightforward.

The procedure to identify seeker and acceptors involved in the round constitutes an architectural trait of the market. Instead of (uniform) random drawing, one could attribute different probabilities to different traders. Seekers could be given different probabilities based on their behavioral state (e.g. one state might entail a higher probability than other states). Once a seeker has been identified (that is, conditionally), acceptors could be given different probabilities based on how their behavioral state matches the one of the seeker (e.g. having the same behavioral state, or a state defined as complementary to the one of the seeker, might entail a higher probability than other states). For example, noise traders, or particular types of them, might be made more likely to meet other noise traders.

Alternatively, acceptors might be given different (conditional) probabilities based on some measures of "closeness" to the seeker, in ways not related to behavioral states. These measures could be proxies for diverse aspects, ranging from sheer size (and hence "visibility") of the acceptor, to spatial closeness, to "institutional" closeness. Finally, one could eventually model mechanisms of reputation and market loyalty, so that with high probability a seeker samples acceptors that have a good "public" reputation and/or with whom he has successfully dealt in the past. This requires the introduction of information on traders' failures (see remarks at the end of Section 2.4).

In general, $1 \leq k \ll \text{card}(T' \in \mathcal{T} : f^\alpha[T'] = 1)$ forces trading interactions to be non-global, to an extent measured by $k/\text{card}(T' \in \mathcal{T} : f^\alpha[T'] = 1)$. If, in the two-stage procedure, acceptors selection (second stage) depends on the seeker (first stage), non-globality can be interpreted as locality in terms of some measure of "closeness".

Also λ is an architectural parameter: it expresses the relative degree of “power” of seeker and acceptor in forming the price of a transaction. If $\lambda = 1$, the transaction price will coincide with the seeker’s reference price, while if $\lambda = 0$ it will coincide with the acceptor’s reference price. The procedure to determine prices, as well as that to determine completion schedulings, are architectural traits of the market, too. As we mentioned already, the experimenter could specify them in a different fashion. In particular, combining them with asymmetric and possibly diversified censoring of trading documents (i.e. what is disclosed upon meeting depends on the trader’s role in the encounter, and possibly on the type of trader), one could attempt to model “order-driven” markets as distinguished from “price-driven” markets as the ones described here (for a discussion of different market types, see Tordjman, 1998).

2.2 Completing transactions

When the board transaction counter shows N , there is a (possibly empty) set of outstanding transactions which are in completion phase for the buyer

$$\sum_{j=1}^{h[o]+dh_1[o]-1} \nu_j \leq N, \quad c_1[o] = 0$$

and/or for the seller

$$\sum_{j=1}^{h[o]+dh_2[o]-1} \nu_j \leq N, \quad c_2[o] = 0$$

Correspondingly, payments and deliveries occur *up to the current cash and asset availability of the traders involved*, according to a pseudo-simultaneous procedure. By *pseudo-simultaneity* we mean, loosely speaking, that failure in some of the scheduled payments (deliveries) depends solely on actual cash (asset) shortages on the side of the traders involved, and not on the ordering in which payments (deliveries) are performed ¹¹.

At the end of the procedure, all traders involved will have new levels of cash and asset, and updated completion flags for transactions in their note-pads. Transactions that have been completed on both sides become inactive. Nevertheless, they are not removed from \mathcal{O} , and their records are not removed from traders’ note-pads. In fact, the information in them might still be useful. One could interpret this by saying that transaction history is fully retained in the system. Of course, a dissipation mechanism could be contemplated at the system level and/or within traders’ note-pads: a certain number of minutes after conclusion on both sides, one could remove transactions from the system, their prices from the board tape, and their records from traders’ note-pads.

Last, the new levels of asset of traders making or receiving deliveries are checked to determine who must be added to, or removed from, the disclosure sheet in the board display.

¹¹Here is an example: two payments are outstanding at the same moment: T has 10 dollars in his pocket, and owes 5 to T' ; T' has 2 dollars in his pocket, and owes 4 to T'' . If the first payment is considered first, then both payments will be performed. On the other hand, if the second payment is considered first, it will fail although T' can actually count on $5 + 2 = 7 > 4$ dollars. In order to make the procedure pseudo-simultaneous, one must find an ordering of outstanding payments (deliveries) which avoids situations like the one described above. We devised an algorithm for doing this, which is described in detail in Bertè (1998).

2.3 Updating the trading documents: Targeting/Positioning

When the board transaction counter shows N , there is a set of traders that:

- have just concluded a transaction, and/or
- are at the end of a minute on their watches

$$\sum_{j=1}^{H[T]} \nu_j[T] = N$$

The set is certainly not empty, as it always contains at least the seeker and the acceptor who have concluded the transaction bringing the board transaction counter from $N - 1$ to N . What is important to notice is that other traders (who have not just concluded a transaction) might be in the set as well because of time passing by on their watches.

Each trader in this set updates his acceptor and seeker sheets (targeting), and consequently repositions himself in the room. He does so using the targeting/positioning algorithms in his manual (which might differ depending on his behavioral state).

2.4 Expulsion: leaving the room

Suppose T is the buyer in a given transaction o ($b[o] = T$). If he has not performed his payment $p[o]$ by the end of minute $h[o] + dh_1[o]$ ¹², he is allowed to extend the (buyer) completion phase by one minute using one bonus. Similarly, T can extend the (seller) completion phase of a transaction in which $s[o] = T$. Bonus-minutes can be used in sequence and in parallel; that is, to extend the completion phase of one transaction several times, and/or the completion phases of several transactions simultaneously.

At the end of each minute H on the board clock¹³, each trader will request a certain number of bonuses, say $dB[T] \geq 0$, to extend completion phases to the next minute. If $B[T] \geq dB[T]$, the bonus-minutes are awarded and used. The counter is updated correspondingly: $B[T] \leftarrow B[T] - dB[T]$. On the other hand, if $B[T] < dB[T]$, the bonus-minutes are not awarded and the trader's expulsion flag $ex[T]$ is switched from 0 to 1.

All concluded transactions involving T for which neither payments nor deliveries have occurred yet (both completion flags = 0), are simply "canceled". Technically, their completion flags are switched to 1 in traders' note-pads as if they had been completed, although payments and deliveries associated to them will never be performed. A sort of *bankruptcy procedure* is then implemented.

Suppose $ex[T] = 1$ following a failed payment in the amount of $p[o]$. T still might have $0 < m[T] < p[o]$ in cash, and $q[T] > 0$ in asset units. Conversely, if $ex[T] = 1$ following a failed delivery, T still might have $m[T] > 0$ in cash, and will necessarily have $q[T] = 0$ ¹⁴. Residual cash and asset units, if any, will be distributed to complete transactions concluded by T for which the other trader has already performed his payment or delivery (T 's side completion flag = 0, while the other completion flag = 1). With his residual $q[T]$, T can cover deliveries for

¹²This coincides with the board transaction counter showing $\sum_{j=1}^{h[o]+dh_1[o]} \nu_j$.

¹³This corresponds to the board transaction counter showing $\sum_{j=1}^H \nu_j$.

¹⁴This is because the failed delivery is bound to concern one unit of asset.

payments he has already received, or give back units he has already taken but not yet paid for. Similarly, with his $m[T]$, T can cover payments for deliveries he has already received, or give back cash he has already taken but not yet delivered for.

Regarding the asset, one can perform up to $q[T]$ deliveries/units restitutions, and we use a chronological ordering based on completion scheduling. Regarding cash, one can perform one or more payments/cash restitutions whose global amount does not exceed $m[T]$, and we use an increasing price order ¹⁵.

Likewise a regular completion round, the bankruptcy procedure modifies the levels of cash and asset of the traders involved, including the one under bankruptcy. T 's side completion flags for the transactions that are covered through the bankruptcy procedure are switched to 1 in the traders' note-pads. Finally, transactions already concluded by the other party which did not get covered on T 's side through the bankruptcy procedure are "canceled" as well (T 's side completion flag is switched to 1), with a net loss (as usual, inactive transactions are not removed from \mathcal{O} , and their records are not removed from traders' note-pads).

T is then removed from \mathcal{T} ; he leaves the room (through the "door") irreversibly. Unlike records relative to completed (inactive) transactions, records relative to expelled (irreversibly inactive) traders are removed from other traders' note-pads.

A delicate issue is that of a trader's *residual asset and/or cash endowments* (if any) upon *expulsion*. In fact, even after the bankruptcy procedure, T might have $m[T]$ and/or $q[T] > 0$. Letting T walk out of the room with them, i.e. eliminating the remaining endowments, would create outflows of cash and/or asset units from the room. Cash outflows are conceivable, but units outflows might be troubling. In particular, this is the case whenever FTR is used to represent trading in a stock, whose overall number of units (shares) ought to remain constant ¹⁶. An easy solution to the problem is to pool the residual units of expelled traders into a "fund" from which traders entering the room draw their initial asset endowments (see below).

Another remark is in order here. Suppose only some particular traders were allowed behind windows. Then, the expulsion of all such traders would automatically annihilate the whole trading system, as no one could accept transactions anymore. Similarly, in the case in which only some particular traders were allowed in chairs, the system would collapse for lack of seekers if all those trader were expelled. The "irreversible" collapse of a market is a rare but possible event, which could be produced by FTR though this route. Relatedly, "non-irreversible" market collapses could be produced if all traders allowed behind windows (in chairs) chose to stand by the third wall (i.e. to temporarily renounce involvement).

The initial number of bonus-minutes $B_{max} \in \mathbb{N}^I$ is another architectural parameter of the market. Large values of B_{max} increase the likelihood that concluded transactions will eventually be completed, and decrease the likelihood of traders eventually being expelled from the room. On the other hand, large values of B_{max} allow for substantial departures from the agreed upon completion schedulings, weakening the role of the latter as both a "decision variable" and a "disciplining device".

¹⁵This maximizes the *number* of payments/cash restitutions one covers, but is an arbitrary choice. Other orderings could be implemented by the experimenter: for example, a decreasing price order would give priority to larger payments/cash restitutions. Of course, a further alternative would be a random order: payments/cash restitutions would be listed in random order, and scanned until one (if any) is found which does not exceed $m[T]$; this item would be covered and eliminated from the list, and $m[T]$ would be reduced accordingly. The scanning would resume until a payment/cash restitution is found (if any) which does not exceed the new $m[T]$, etc.

¹⁶Unless, in further developments, one attempts to model share issues, buy-backs, etc.

Notice that, for the time being, we do not maintain information relative to traders' failures anywhere in the system. The rationale is that the market is one in which honoring of agreements is institutionally enforced through an expulsion penalty (which is stronger the smaller B_{max}). Moreover, as we will see in Section 3.1, traders can be assumed to be fully aware of this, and to behave accordingly. Thus, a memory of "bankruptcy reputation" is unnecessary.

Obviously, one might want to modify this: information relative to traders' failures could be introduced at several levels. Identities and $B[T]$'s of all traders, or of all traders whose number of available bonus minutes is below a certain threshold, could be posted on the board disclosure sheet and updated in "real-time"; in this case, the information would be complete, never obsolete, and public. Alternatively, traders might be required to report their $B[T]$'s on their acceptor and seeker sheet (and not to censor it); in this case, the information would be complete, but passed and registered in note-pads only upon pair-wise encounters. In both cases, the flows of information relative to available bonus-minutes is institutionally regulated via the board or the trading documents. The second scenario limits flows to pair-wise encounters: to know about T' , T must meet him. Moreover, the information T keeps in store will be relative to the time in which the last encounter took place, and thus subject to obsolescence (until T meets T' again). Last, one might add a slot to a trader's note-pad records of other traders. Besides trading documents, T might keep counts of delay-minutes on payments and deliveries inflicted on T himself by each T' he has dealt with. In this case, the flows of information are not institutionally regulated; they can obviously be obsolete, and are incomplete, as T will count only those delays that affected him directly¹⁷.

If information on traders' failures is introduced in any of the above ways, it could obviously be used in decision-making processes, i.e. become one of the inputs of the positioning/targeting and transaction-selection algorithms. Besides selection among transactions made available by sampled acceptors, information on traders' failures could also be used to orient the seeker's sampling of acceptors in a trading round.

Finally, note that with some easy additions to the current version of FTR, some traders (i.e. "market makers") could be allowed to access outside credit (at least up to a ceiling) rather than use bonus-minutes. Relatedly, those traders would perform as a sort of "clearing house" for the market.

2.5 Entering the room

Expulsion constitutes a natural *death process* for the system. A *birth process* could be introduced as well¹⁸. Through births one can represent inflow of new investors, and appearance of *new types* of traders; that is, of traders characterized by novel behavioral and cognitive patterns. As death by expulsion, birth could be anchored to the board (transaction-based) clock by admitting traders into the room at the beginning of each minute.

¹⁷A further alternative would be to have information on $B[T]$'s pass through *costly phone lines*. This hints to a whole other class of information flows that could be introduced in the model; that is, information regarding the trading process which can be accessed possibly by a subgroup of traders, and possibly at a cost. We have taken phones as a metaphor for information from "outside the room", but one could introduce a second row of phones conveying information from "inside the room" to traders who know the required phone numbers –access codes–, with a given fee. See the discussion in Milgrom, North and Weingast (1990).

¹⁸Technically, the set \mathcal{O} is augmented by transaction conclusion. Transactions become then inactive over time as they get completed, but as we have seen they are never removed from \mathcal{O} . A birth process augments \mathcal{T} : traders become then inactive over time if and only if they are expelled, and as we have seen they are removed from \mathcal{T} .

When generating new traders, the experimenter has to decide, among other things, how to *initialize their cash and asset endowments, and their manuals* (i.e. cognitive and behavioral patterns). Cash and asset endowments of entrants create inflows of cash and units in the room.

Entry can be formalized as a two-stage random procedure. At the beginning of a minute, the number of entrants $n_{E(H)}$ is drawn from a distribution \mathcal{N}_E on \mathbb{N}^I . Draws are independent across minutes. Of course, one could *endogenize* \mathcal{N}_E . For instance, one could allow the number of entrants to depend, in probability, on past average returns on the market.

The units in the “fund” obtained by pooling residual asset endowments of expelled traders, if any, are evenly distributed among new entrants. Then, cash levels and manuals are initialized. This is done through $n_{E(H)}$ independent draws from a distribution, say $\mathcal{I}_{(H)}$, on $\mathbb{R}_+^I \times$ the (current) “space of manuals”. Clearly, the nature of the latter will depend on how the algorithms in the manual are formalized, and might be quite complex.

One way of specifying $\mathcal{I}_{(H)}$ which works regardless of the nature of the space of manuals, is *cloning*; that is, drawing at random from a weighted group of traders (i.e. allowing for different probabilities for the various members), and attributing to the entrant under consideration the cash level and manual of the selected trader. Obviously, one can use a noisy version of cloning, superimposing an (independent) error to the cash level and manual of the trader being cloned. This brings back issues related to the nature of the space of manuals, though, as one must define an error on such space ¹⁹.

Also, one has to specify the group of traders from which to draw. For some experiments, one might want to use the group of *incumbents* (i.e. the traders already in the room at the end of $H - 1$) weighted by wealth (say $m[T] + p_{(N)}q[T]$, where cash and asset levels are the ones immediately after $N = \sum_{j=1}^{H-1} \nu_j - 1$ –the last transaction of minute $H - 1$). Of course, other measurements of performance (e.g. realized returns) over a certain historical record might be used to weight traders.

Some words cautions are in order: weighting incumbents by wealth or other performance measurements, one makes the implicit assumption that the latter are known to potential entrants who, although with a random element at play, target the most “successful” incumbents. Moreover, attributing to a new entrant the manual of an incumbent, even with an error superimposed to it, one implicitly assumes that there is very little appropriability to the behavioral and cognitive patterns embodied in manuals. This becomes critical when incumbents are allowed to learn, i.e. to evolve such patterns. Similar considerations apply for cash endowments: one makes the implicit assumption that the new entrant will be capable of starting out at the same cash-size of the cloned incumbent, possibly with a noise.

For other experiments, one might want to use the group of incumbents, with weights specified in terms of types of traders. When cloning, such weights are proxies for the tacitness and appropriability of the behavioral/cognitive patterns characterizing each incumbent type (high weights corresponding to low tacitness and appropriability).

In some cases, one can construct an *ad-hoc* weighted group of pseudo-traders from which to draw. This will give manuals and cash endowments distributions for entrants, tailored to the experiment. In particular, this is the way to go when using entry to experiment with immission of specific *new* types of traders in the system (weights will correspond to probabilities for each

¹⁹For example, one might copy strings expressing algorithms, with non-zero probability of mistakes in various positions along the strings themselves. This would be a sort of population-version of the exploration process modeled at the individual level by Marengo and Tordjman (1996).

new type).

In full generality, entry allows experiments on the invadability of particular “ecologies” of behavioral/cognitive patterns, with their population-level dynamics of replication, by novel patterns that enter the market along its unfolding history.

2.6 The “external value(s)”

As already mentioned, we assume the existence of one (or more) external variable(s) affecting the the “real” value of the object which the financial asset denominates (i.e. so called fundamentals). We considered one such variable, labeled “external value”, for simplicity, and assumed it to evolve on the minute-scale following the board clock. How one specifies the exogenous evolution of $Z_{(H)}$ is obviously crucial to simulation experiments whenever a non-negligible number of traders uses information from the phones. As a first approximation we set, at the beginning of each minute H

$$Z_{(H)} = z_{(H)} + \zeta_{(H)}$$

where $z_{(H)}$ is a *systematic component*, and the $\zeta_{(j)}$, $j = 1, 2, \dots$ are independent draws from a probability distribution \mathcal{Z} (usually, but not necessarily, a 0-mean normal). Obviously the variance can be set to 0, reducing $Z_{(H)}$ to coincide with the systematic component.

Notice that we are allowing for two levels of randomness relative to the external value: First, $Z_{(H)}$ can itself contain a *fluctuation about the systematic component*. Second, traders who access the external return might superimpose to the fluctuation an independent *reading error* ($e[T]$). Keeping these two levels of randomness separate increases the flexibility of FTR in representing various scenarios.

A dynamics must be specified for $z_{(H)}$: in certain settings one might want to keep $z_{(H)}$ constant, but any non-constant exogenous “trend” can be implemented. Moreover, one can implement a *random walk* using an endogenous formulation for the systematic component: $z_{(H)} = Z_{(H-1)}$ ²⁰. A non-0-mean for the i.i.d. steps $\zeta_{(j)}$, $j = 1, 2, \dots$ would then represent a *drift*.

In the experiments described in Chiaromonte and Bertè (1998), $z_{(H)}$ is constant, $var(\zeta_{(H)}) = 0$, and traders have non-zero variance reading errors.

3 Behavioral repertoires and market features

The algorithms in a trader’s manual can use as input any variable on the board, from the phones, or “internal” to the trader himself. In particular, they can use (public) information from the board display, any (outside) information obtained through the phones, and any (private) information recorded in the trader’s note-pad. In full generality, we only assume the targeting algorithms to be such that

$$p_b^\alpha[T] \leq p_s^\alpha[T] \text{ , } p_b^\sigma[T] \leq p_s^\sigma[T]$$

This is a consistency requirement on the mechanism forming reference prices (see trading documents in Section 1).

²⁰Rigorously, this means setting $Z_{(H)}|Z_{(H-1)} \sim Z_{(H-1)} + \zeta_{(H)}$. Notice one has to specify a initial value or distribution for $Z_{(0)}$.

It is important to stress that the model can host any kind of *behavioral heterogeneity*. Multiple behavioral/cognitive patterns coexisting within the “brain” (i.e. the manual) of a single traders (as for example in Marengo and Tordjman, 1996) can be represented by diversifying algorithms across states $r[T]$. Behavioral and cognitive diversity among traders can be represented by diversifying algorithms in corresponding states, or introducing different states, across traders. In all cases, heterogeneity can be simply in terms of parameter values, or extend to form and nature of the algorithms themselves.

While some of the algorithms embody the cognitive/behavioral repertoire that traders are “free” to chose and evolve, other algorithms might embody *internal representations* of the rules which particular institutional architectures impose upon “orderly” trading behavior. Let us consider a few examples.

3.1 The honoring constraint

The market is one in which agreements are supposed to be honored. As we have seen, traders who fail to complete transactions they concluded according to their scheduling are subject to strong penalties (expulsion after B_{max} bonus-minutes). We assume traders to be aware of this, and behave accordingly. We have also discussed a corollary to this assumption: since missed completion ought to be an exception, information on traders’ failures is not maintained anywhere in FTR (and is not used in traders’ decision-making).

Still, completion of outstanding transactions might not be the main factor determining traders’ decisions when actual exchange events are not on the immediate horizon. In order to capture this, we assume all traders to have $R[T] \geq 2$ behavioral states, the first of which is labeled *red-alert*. This state implies the use of targeting/positioning and transaction-selection algorithms particularly aimed at collecting the cash and asset units necessary to complete outstanding transactions. We then take a trader to be in red-alert ($r[T] = 1$) if and only if he is in completion phase (as a buyer and/or a seller) for at least one of his outstanding transactions.

If red-alert algorithms are effective, and minutes (on the board clock) long enough on average in terms of transactions, traders ought to be able to complete almost all the transactions they concluded according to their scheduling.

Notice that we link the red-alert state to having transactions in completion phase, and not to knowing that $B[T]$ is approaching 0. In other words, we assume traders not to willingly take advantage of the existence of bonus-minutes. This can be considered a conservative (and possibly sub-optimal) behavior if traders are aware of the existence of the bonuses, and know the exact value of B_{max} . On the other hand, it can be easily interpreted as a reasonable (if not mathematically optimal) rule, if traders attribute a large enough dis-utility to being irreversibly expelled from the room. Finally, it can be taken as an institutionally shaped “ethical” trait.

3.2 Synchronous vs asynchronous completion markets

The market might be one in which trading goes on without interruption. In this case it would not make sense to constrain the completion phases of non-spot-spot transactions to some particular time intervals. On the other hand, the market might be one in which *trading is divided into discrete periods*, say *days*, and payments and deliveries associated to transactions concluded during one day are supposed to take place within the day itself. This bounds all transactions concluded during the last minute of each day to be spot-spot. Once more, we assume the traders

to be aware of this, and to behave accordingly.

Going one step further, we can then assume that in such a setting traders will schedule all exchanges that are not spot to occur during the last minute. Thus, the completion phases (for both buyers and sellers) of all non-spot-spot transactions concluded during a given day will concentrate in the last minute of the day itself. Correspondingly, traders will be in red-alert only then, if ever. This further divides each trading day into what we could call a *speculation* or *arbitrage period* (all minutes before the last), and a *completion rush* (the last minute) characterized by spot-spot trading.

This description is accurate with the exception of traders extending completion phases through bonus-minutes: delayed payments and deliveries will carry over to the first minute(s) of the following trading day. Hence, what we named the speculation period of a day could still see some completion phases and some traders in red-alert. If red-alert algorithms are reasonably effective, last minutes are long enough on average in terms of transactions, and traders are not given too many bonus-minutes, one should still notice a fairly clear overall differentiation between trading all along the day, and trading in the last minute. If not, one might actually have a three-period day with a completion rush in the initial minute(s), a speculation phase, and the final completion rush –whose “tails” generate completion rush in the initial minute(s) of the following day, etc.

In symbols, take $H^* \in \mathbb{N}^I \setminus \{0\}$ to be the length in minutes of a trading day (which is assumed to be the same for all days). The targeting/positioning algorithms of each T will be required to be such that

$$D_b^\alpha[T], D_s^\alpha[T], D_b^\alpha[T], D_s^\alpha[T] \subseteq \{(0, mH^* - H)\} \times \{(0, mH^* - H)\}$$

under all behavioral states. H is as usual the current minute on the board clock, and $m = 1, 2, \dots$ is a day counter.

3.3 Imitation

An interesting issue concerns the representation of imitation phenomena. Imitation is constrained by what traders know about each other through the board disclosure sheet, and the records they keep in their note-pads. Still, even without introducing any further information items in these loci, the targeting/positioning algorithms of a trader T could use as input the information on size from the board display, combined with what he can infer on other traders' targeting from the records he maintains of their acceptor and seeker sheets. Let us make a simple example: T could take up the flags and the reference prices of T' as they were the last time the two met (or more generally use them in his updatings), if he knows from the mandatory public disclosure that $q[T']$ has exceeded a given level $\xi[T'] (\geq Q)$, and the meeting occurred less than $\eta[T']$ minutes ago. In essence, this pattern embodies something like: “if you have recent information on what George Soros did, do the same” (or more generally, behave accordingly).

A certain number of traders in the room could then be characterized by this imitative targeting behavior (possibly with different parameters $\xi[T]$ and $\eta[T]$) whenever not in red-alert. More generally, a behavioral state (say $r[T] = 2$) implying such behavior could be introduced for all traders, together with conditions under which the state is entered and exited.

3.4 (Functional) differentiation of traders

One could also qualify traders' algorithms, and hence use what we label behavioral states, to define *different roles* in the trading process. We may call *investor* a trader whose targeting/positioning algorithms are such that

$$f_b^\alpha[T] = f_s^\alpha[T] = 0 \text{ and therefore } f^\alpha[T] = 0$$

all along. An investor is a trader who is not allowed behind windows, is aware of it, and behaves accordingly. On the other hand, we may call *market maker* a trader whose targeting/positioning algorithms allow for the acceptor-side flags to be equal to 1. A market maker will sometimes, or even always, be behind a window.

Notice that investors, who metaphorically correspond to “real-economy” operators (e.g. individual savers, commodity-producing firms, etc.), will never be in two places simultaneously. Only market makers, who metaphorically correspond to financial operators, can be at a window and in a chair at the same time.

Targeting/positioning algorithms, as well as selection ones, will be further differentiated to embody differences in aims and approaches to trade between the two types of traders.

Once more, a certain number of traders in the room could then be characterized as market makers, and the remaining as investors. Notice that, as investors can only be seekers, while market makers can be both seekers and acceptors, a necessary condition for the system to function is the existence of at least one market maker.

3.5 A tentative behavioral taxonomy for spot-spot trading based on “price assessment”

Assume that traders limit themselves to transactions that are spot on both sides. Whenever updating, they set (completion scheduling options)

$$D_b^\sigma[T] = D_s^\sigma[T] = D_b^\alpha[T] = D_s^\alpha[T] = \{(0, 0)\}$$

and targeting/positioning, as well as transaction-selection (by seekers), concern only flags and reference prices.

Since completion scheduling (and thus, possible information on traders' failures) plays no role, we attribute to all traders –who do seek– a transaction-selection algorithm based on mere price ranking among available transactions in a trading round ²¹.

As far as targeting/positioning is concerned, flags and/or reference prices are assumed to be set in relation to some “peg”, which implicitly embodies the expectations of the trader on the value of the asset. For example, “fundamentalists” will form their peg based on information on the fundamental variable(s), while “chartists” will attempt to detect a structure in the past price dynamics, etc. We call the peg a *price assessment*. In symbols, let

$$G_\theta(x_o, x) = \max\{ 0, x_o + g_\theta(x - x_o) \}$$

be a function from \mathbb{R}_+^2 to \mathbb{R}_+^1 parameterized through a vector θ . $g_\theta(\cdot)$ is from \mathbb{R}^1 to \mathbb{R}^1 ; in particular, we will consider

$$g_\theta(y) = \theta_1 y \text{ Ind}(|y| \geq \theta_2)$$

²¹Notice that ranking reference prices of acceptors or transaction prices is equivalent, as the latter are just convex combinations (with the same λ) between the former and the seeker's reference prices.

where $Ind(\cdot)$ is the indicator function of the argument condition, $\theta_1 \in \mathbb{R}^l$ and $\theta_2 \in \mathbb{R}_+^l$. This represents a linear form which is flat at 0 in a neighborhood of 0 of radius θ_2 , and has slope θ_1 outside it.

All traders use $G_\theta(\cdot, \cdot)$ to form their *price assessment*. x_o can be interpreted as a *center*: within a θ_2 -neighborhood of x_o , there's no reactivity to the difference $x - x_o$, and the price assessment is set at the center itself, while outside the neighborhood, one has a linear reaction whose sign and size are expressed by the sign and size of θ_1 .

Also, all traders update flags and reference prices based on their price assessment. However, we differentiate traders in three respects, namely

- what variables x_o, x enter the computation of the price assessment (*what information is used*)
- how the price assessment components x_o and $g_\theta(x - x_o)$ are employed in updating flags and reference prices (*how such information is used for targeting/positioning*)
- the values of the parameters for the price assessment computation ($\theta_1[T], \theta_2[T]$), and of any parameters for the updating of flags and reference prices (see below)

In terms of **variables entering the price assessment computation**, one could list at least five possibilities:

1. A **(strong) fundamentalist** trader with access to a noisy version of the *current external value*. For instance, assume that the asset is a stock. Further, for simplicity of exposition, suppose that the external value $Z_{(H)}$ is not the return of the related firm in H , but already the “equilibrium” capitalization (whatever that might mean...) of the whole flow of present and future returns. The fundamentalist will simply take $Z_{(H)} + e[T]$ as price assessment. This corresponds to setting

$$x_o = x = Z_{(H)} + e[T]$$

Clearly $g_{\theta[T]}(x - x_o) = 0$, so the parameters $\theta_1[T], \theta_2[T]$ are irrelevant in this case.

2. A **(quasi) fundamentalist** trader with access only to the *history of external values*, who extrapolates the current external value with a moving average, and then computes his price assessment setting

$$x_o = Z_{(H-1)} \quad , \quad x = \sum_{j=0,1,\dots} \chi_j[T] Z_{(H-1-j)}$$

(the parameters $\chi_j[T]$ are non-negative weights adding up to 1).

3. An **adaptive trader** or **chartist** (see Brock et al. 1992), looking at the time series of *prices* (which is public information, up to the very last transaction), who extrapolates the next transaction price with a moving average, and then computes his price assessment setting

$$x_o = p_{(N)} \quad , \quad x = \sum_{j=0,1,\dots} \psi_j[T] p_{(N-j)}$$

(the parameters $\psi_j[T]$ are non-negative weights adding up to 1).

4. As a special case when $\psi_0[T] = 1$, $\psi_j[T] = 0, j \neq 0$, we obtain a **noise trader** (see Grossman and Stiglitz, 1980) that, somewhat like the fundamentalist in (1) with respect to $Z_{(H)} + e[T]$, simply takes as price assessment the *last transaction price* $p_{(N)}$:

$$x_o = x = p_{(N)}$$

Again $g_{\theta[T]}(x - x_o) = 0$, so that $\theta_1[T], \theta_2[T]$ are irrelevant.

5. A **sun-spotter** who, in the simplest version, takes as price assessment a *random variable* V *completely unrelated to both trade and the “real” value of the object denominated by the asset*. This corresponds to setting

$$x_o = x = V$$

Once more $g_{\theta[T]}(x - x_o) = 0$, so that $\theta_1[T], \theta_2[T]$ are irrelevant. In a somewhat different version, a sun-spotter could use the current price as center, and *react* to a random variable W (unrelated to trade and “real” value) through the g_{θ} -formulation ²²:

$$x_o = p_{(N)} \quad , \quad x = W$$

6. An **imitator** who chooses *another trader* T' *as target* (based for example on size information from the board disclosure sheet, and information on last encounters from his records), and then computes his price assessment setting

$$x_o = p_{(N)} \quad , \quad x = \hat{p}_{(\cdot)}^{(\cdot)}(T')[T]$$

Needless to say, one might define, and experiment with, traders characterized by much more sophisticated inferential procedures (they could be skilled econometricians, use neural nets, etc.).

In terms of **how the price assessment components are employed in updating flags and reference prices**, consider as illustration the following examples:

1. A **“Take-Action”** trader who updates his seeker flags on the basis of $g_{\theta[T]}(x - x_o)$ as:

$$\begin{aligned} g_{\theta[T]}(x - x_o) > 0 &\implies f_b^{\sigma}[T] = 1 \quad , \quad f_s^{\sigma}[T] = 0 && \text{seek to buy} \\ g_{\theta[T]}(x - x_o) = 0 &\implies f_b^{\sigma}[T] = 0 \quad , \quad f_s^{\sigma}[T] = 0 && \text{hold} \\ g_{\theta[T]}(x - x_o) < 0 &\implies f_b^{\sigma}[T] = 0 \quad , \quad f_s^{\sigma}[T] = \text{Ind}(q[T] \geq 1) && \text{seek to sell} \end{aligned}$$

Notice that whether buying (selling) is associated to a positive (negative) $x - x_o$ or vice-versa, depends on the sign of $\theta_1[T]$; both links can be considered and represented. The indicator function in the last line makes seeking to sell conditional to having at least one asset unit in store ²³.

Seeking reference prices, on the other hand, do not depend on the trader’s price assessment: he will pursue the action he selected to the limits of his current cash endowment ²⁴ setting

$$p_b^{\sigma}[T] = m[T] \quad , \quad p_s^{\sigma}[T] = 0$$

so if $f_b^{\sigma}[T] = 1$, T is willing to buy at any price within his budget constraint, while if $f_s^{\sigma}[T] = 1$, he is willing to sell no matter how low the price. Since Take-Action embodies

²²one could take for example $W \sim N(p_{(N)}, \sigma^2)$, so that $(x - x_o) \sim N(0, \sigma^2)$.

²³Recall trading is spot-spot, so deliveries cannot be postponed.

²⁴Likewise deliveries, payments cannot be postponed.

such an extreme approach, we confine it to seeking and assume a Take-Action trader never to be an acceptor ²⁵:

$$f_b^\alpha[T] = 0 \quad , \quad f_s^\alpha[T] = 0$$

Consequently, acceptor reference prices are irrelevant.

2. A **“Form-a-Spread”** trader is always available to exchange, within the limits of his asset endowment, and of a *share* of his cash endowment. Moreover, he can both seek and accept transactions, and does not diversify his flags and reference prices in the two roles. Thus, he sets

$$f_b^\sigma[T] = f_b^\alpha[T] = 1 \quad , \quad f_s^\sigma[T] = f_s^\alpha[T] = \text{Ind}(q[T] \geq 1)$$

The price assessment components x_o and $g_{\theta[T]}(x - x_o)$ are used to form reference prices, as to define a *spread*. The calculation involves both a spread parameter $\varepsilon[T] \in \mathbb{R}_+^1$, and a caution parameter $\gamma[T] \in [0, 1]$:

$$\begin{aligned} g_{\theta[T]}(x - x_o) > 0 &\implies p_b^\sigma[T] = p_b^\alpha[T] = \min\{x_o, \gamma[T]m[T]\} \\ &\quad p_s^\sigma[T] = p_s^\alpha[T] = x_o + 2\varepsilon[T] \\ g_{\theta[T]}(x - x_o) = 0 &\implies p_b^\sigma[T] = p_b^\alpha[T] = \min\{x_o - \varepsilon[T], \gamma[T]m[T]\} \\ &\quad p_s^\sigma[T] = p_s^\alpha[T] = x_o + \varepsilon[T] \\ g_{\theta[T]}(x - x_o) < 0 &\implies p_b^\sigma[T] = p_b^\alpha[T] = \min\{x_o - 2\varepsilon[T], \gamma[T]m[T]\} \\ &\quad p_s^\sigma[T] = p_s^\alpha[T] = x_o \end{aligned}$$

Again, the sign implications depend on the sign of $\theta_1[T]$. This can be interpreted as follows: T considers his price assessment $x_o + g_{\theta[T]}(x - x_o)$ as relative to some generic future, but not entirely reliable as a point-evaluation. Thus, he does not use $x_o + g_{\theta[T]}(x - x_o)$ to form reference prices for the very next transaction he will engage in. Instead, T uses the sign of $g_{\theta[T]}(x - x_o)$, which he trusts to be reliable, to orient upwards or downwards a spread anchored to the center x_o .

Note that in all cases in which one sets $x_o = x$ (e.g. the first type of fundamentalist and the noise trader in the previous classification), the price assessment reduces to x_o and the spread is always symmetric about it. The spread can then be interpreted as some sort of interval-evaluation, and $\varepsilon[T]$ as a measure of the uncertainty the trader attributes to the price assessment pivoting it (see Chiaromonte and Bertè, 1998).

Relatedly, note also that the experimenter could turn $\varepsilon[T]$ into an endogenous variable ²⁶. In analogy to what frequently suggested in the literature, $\varepsilon[T]$ could depend, for example, on market volatility and other variables.

3. Premises and flags are the same as for “Form-a-Spread”, but a **“Form-a-Divide”** trader forms a unique reference price acting as a divide between buying and selling (a caution parameter $\gamma[T] \in [0, 1]$ is again involved in the calculation):

$$\begin{aligned} p_s^\sigma[T] = p_s^\alpha[T] &= p^* = x_o + g_{\theta[T]}(x - x_o) \\ p_b^\sigma[T] = p_b^\alpha[T] &= \min\{p^*, \gamma[T]m[T]\} \end{aligned}$$

²⁵In the role of seeker, a Take-Action trader will be the party selecting a transaction among the ones made available by the sampled acceptors. Thus, although he is willing to buy/sell at very extreme prices, he will at least chose the most convenient price.

²⁶Given the modularity of the computer-implementation of FTR, and the underlying structure of the LSD platform, turning parameters into variables, and specifying their evolution as a function of other variables and parameters in the system, is a straightforward exercise (see Bertè, 1998, and Valente, 1997 –concerning the LSD platform).

This can be interpreted as considering the price assessment $x_o + g_{\theta[T]}(x - x_o)$ a reliable point-evaluation for the close future, and hence using it in the very next transaction: any price below the divide $p^* = x_o + g_{\theta[T]}(x - x_o)$ is seen as a buying opportunity (within the limits of a share of $m[T]$), and any price above the divide as a selling opportunity.

In all cases in which one sets $x_o = x$, “Form-a-Divide” is just a special case of “Form-a-spread”, with $\varepsilon[T] = 0$.

Considering some combinations of the two foregoing classifications, the following table presents a rudimentary taxonomy of trader-types.

SOME TRADER-TYPES			
	Take-Action	Form-a-Spread	Form-a-Divide
$Z_{(H)} + e[T]$ Strong Fund.		Uses (noisy) information on the current external value. Always available to buy and sell (equivalently as seeker and acceptor). Cautious about his assessments.	Uses (noisy) information on the current external value. Always available to buy and sell (equivalently as seeker and acceptor). Confident in his assessments.
$Z_{(H-1)}$ $Z_{(H-2)}$ \vdots Quasi Fund.		Uses information on the history of external values. Always available to buy and sell (equivalently as seeker and acceptor). Cautious about his assessments.	
$p_{(N)}$ $p_{(N-1)}$ \vdots Chartist	Uses public information on prices. Seeking only, with a very extreme price strategy.	Uses public information on prices. Always available to buy and sell (equivalently as seeker and acceptor). Cautious about his assessments.	
$p_{(N)}$ Noise Trader	Uses the price of the last transaction. Seeking only, with a very extreme price strategy.	Uses the price of the last transaction. Always available to buy and sell (equivalently as seeker and acceptor). Cautious about his assessments.	
V (or W) Sun-spotter	Uses a variable unrelated to both trade and the “real” value. Seeking only, with a very extreme price strategy.	Uses a variable unrelated to both trade and the “real” value. Always available to buy and sell (equivalently as seeker and acceptor). Cautious about his assessments.	Uses a variable unrelated to both trade and the “real” value. Always available to buy and sell (equivalently as seeker and acceptor). Confident in his assessments.

When implementing a sub-set of this taxonomy, the behavioral state $r[T]$ has been used for one classification, and a further indicator (*type*[T]) for the other (see Bertè, 1998²⁷). In particular, the experiments described in Chiaromonte and Bertè (1998) concern a trading room inhabited by Strong Fundamentalists and Noise Traders forming spreads.

²⁷The labeling across the taxonomy is slightly different.

3.6 Evolution of behaviors

Clearly, even if algorithms in traders' manuals are not changing over time, the range and relative weight of behavioral and cognitive patterns at the system-level changes as a consequence of birth, trade (accumulation/decumulation of wealth), and death.

In this sense, FTR as described so far permits the representation of those special cases of evolutionary dynamics driven exclusively by *selection*, and possibly a form of exogenous introduction of novelty.

However, as we have already mentioned, one may introduce a dynamics on algorithms in the manual at the level of each single trader; that is, *learning* (notice the difference between this notion of learning –i.e. evolution of the behavioral and cognitive patterns, and mere cumulation/updating of information while trading²⁸).

The way learning can be formalized will obviously depend on the formal framework in which one chooses to embed the algorithms. When those are parameterized, the most immediate option is to introduce an updating mechanism on the parameters (*parametric learning*), but much more sophisticated options can be devised and implemented, entailing the evolution of the algorithms themselves.

A possible approach to this is the one adopted in Marengo and Tordjman (1996), and Arthur et al. (1997). Whichever framework one uses²⁹, the general idea is that of endowing a trader with a whole set of alternative algorithms to perform a given task. Note that those algorithms need not (although they could) be the ones associated with different $r[T]$'s; one could have a set of alternative algorithms within each $r[T]$. The trader employs one or the other algorithm based on scores of their past effectiveness (which, of course, must be defined and measured), while a random mechanism enlarges the set –mutations, recombinations, etc.

4 The Statistical Office

Any simulation model, as well as any real world history, produces an overwhelming amount of data potentially suitable for analysis. One must therefore choose what subset of the data to store throughout simulation runs, and how to organize them as output. Our metaphor for this is a *statistical office*, which produces statistics about the trading process. It is important to stress that these statistics are not meant for the traders, although some coincide with information flows in the room, but for “outside observers” (that is; for users performing simulation experiments).

We distinguished two classes of statistics that might be of interest: *Demographic* time series, and *Economic* time series. Because of the way time is represented in FTR, time series can be produced on the transaction-scale and/or on the (board clock) minute-scale³⁰.

What follows is an illustration of the time series that the statistical office produces in the current computer implementation of the model. Obviously, the list is far from exhaustive; any other statistics of interest ought to be placed here (the statistical office is the locus for simulation

²⁸Looking at the board, using the phones, meeting.

²⁹Strings, graphs, trees, etc.

³⁰Moreover, because of how the code for FTR is implemented, variables can also be saved on a third scale, namely, that of transaction rounds (regardless of whether they terminate with the conclusion of a transaction). See Bertè (1998) for more details.

Demographic Time Series	
On minute-scale	No. of births (entries) in H -series
	No. of deaths (expulsions) in H -series
	No. of traders in the room at the beginning of H -series

The three time series can be reported on one plot, which summarizes the demographic dynamics through a simulation run ³³.

Economics Time Series	
On transaction-scale	Price of N -series
	No. of traders willing to buy (seek and/or accept) in the round for N -series
	No. of traders willing to sell (seek and/or accept) in the round for N -series
	No. of traders under size discl. right before N -series
On minute-scale	No. of transactions in H -series
	No. of bonus-minutes used by traders in H -series
(ave. on trans.'s in each min.)	Ave. and st. dev. of the price in H -series
	Ave. and st. dev. of the no. of traders willing to buy in H -series
	Ave. and st. dev. of the no. of traders willing to sell in H -series
	Ave. and st. dev. of the no. of traders under size discl. in H -series

For each of the above, one can plot the average H -series, and a $+/-$ one standard deviation envelope about it. A single plot might also contain more than one enveloped average H -series; for example those for the number of traders willing to buy and to sell. Moreover, N -series and corresponding enveloped average H -series can be superimposed on the same plot.

Besides system-level demographic and economic time series, one might want to produce *micro-data*. For example ³⁴:

Anagraphic Data – For each T	
Minute of birth (entrance)	Minute of death (expulsion)

This allows one to generate traders' age distributions relative to any minute H .

Also, throughout each trader's life-time (permanence in the room), one can produce ³⁵

Micro Time Series – For each T	
On transaction-scale	Reference prices in the round for N -series
	Cash endowment right before N -series
	Asset endowment right before N -series
On min.-scale (ave. on trans.'s in each min.)	Ave. reference prices in H -series
	Ave. cash endowment in H -series
	Ave. asset endowment in H -series

³¹Modifying the code to generate any other statistics is very straightforward. Again, see Bertè (1998) for more details.

³²The No. of births (entries) in H -series coincides with the series of $n_{E(j)}$, $j = 1, 2, \dots$

³³The No. of transactions in H -series coincides with the system converting sequence ν_j , $j = 1, 2, \dots$

³⁴The Minute of death (expulsion), only if prior to the end of the simulation.

³⁵“On min.-scale (ave. on trans.'s in each min.)” refers again to overall transactions, and not the ones in which the trader is involved; recall it is overall transactions counting that defines our time in terms of minutes on the board clock. Obviously, for any of the average series here, one could produce also the corresponding standard deviations.

This allows one to study the dynamics of price targeting at the micro level, as well as to follow shares in cash and asset, on the minute-scale. Moreover, one can generate traders' (average) size distributions in cash and asset relative to any minute H .

Let us stress once more that the foregoing illustrations must be considered just as instances of a rich variety of aggregate and micro statistics which FTR can generate.

5 Conclusions

FTR expands upon earlier “artificial markets”, and attempts to provide a simulation environment whereby individual behavioral/cognitive patterns and learning processes, architectural and institutional traits, and time-embedding of events, can be modularly designed and investigated in terms of emerging dynamic properties of the market—including the fate of operators carrying particular behavioral and cognitive features.

The structure of FTR, as well as the statistical outputs it can generate, allow for an easy matching with empirical micro-structural studies of financial markets. Moreover, in the spirit of inter-theoretical comparisons—somewhat alike those pioneered by Axelrod and colleagues in the field of Game Theory—FTR permits “tournaments” amongst different behavioral micro-foundations, the assessment of performances by different trader types (e.g. in terms of relative wealth and survival), and the analysis of the statistical properties of different “ecologies of behaviors”.

Last but not least, one of the main purposes of FTR is to provide a framework through which experiments cannot only be designed, but replicated, incrementally built upon one another, and thus easily compared, by all interested scholars.

References

- Arthur W.B., Holland J.H., Le Baron B., Palmer R., Taylor P. (1997). Asset Pricing under Endogenous Expectations in an Artificial Stock Market. In Arthur et al. (eds.).
- Arthur W.B., Durlauf S.N., Lane D.A. (1997). *The Economy as an Evolving Complex System II*. Adison Wesley, Reading, MA.
- Beltratti A., Margarita S. (1992). Simulating an Artificial Adaptive Stock Market. Mimeo, Turin University.
- Bertè M. (1998). Technical Description of the Financial “Toy-Room”. I.I.A.S.A. Interim Report ???.
- Blume L., Easley D. (1990). Evolution and Market Behavior. *Journal of Economic Theory* vol.58, 9–40.
- Brock W.A. (1997). Asset Price Behavior in Complex Environments. In Arthur et al. (eds.).
- Brock W.A., Hsieh D., Le Baron B. (1991). *Nonlinear Dynamics, Chaos and Instability: Statistical Theory and Economic Evidence*. MIT Press, Cambridge, MA.
- Brock W.A., Lakonishok J., Le Baron B. (1992). Simple Technical Trading Rules and the Stochastic Properties of Stock Returns. *Journal of Finance* vol.XLVII n.5., 1731–1764.

- Chiaromonte F., Bertè M. (1998). Some Preliminary Experiments with the Financial “Toy-Room”. I.I.A.S.A. Interim Report ???.
- De Long J.B., Bradford J., Schleifer A., Summers L.H., Waldmann R.J. (1990). Positive Feedback Investment Strategies and Destabilizing Rational Speculation. *Journal of Finance* vol.45 n.2, 379–395.
- De Long J.B., Schleifer A., Summers L.H., Waldmann R.J. (1991). The Survival of Noise Traders in Financial Markets. *Journal of Business* vol.64 n.1, 1–18.
- Grossman S.J., Stiglitz J.E. (1976). Information and Competitive Price Systems. *American Economic Association* vol.66 n.2, 248–253.
- Grossman S.J., Stiglitz J.E. (1980). On the Impossibility of Informationally Efficient Markets. *American Economic Review* vol.70 n.3, 393–408.
- Goodhart C.A.E., Figliuoli L.(1991). Every Minute Counts in Financial Markets. *Journal of International Money and Finance* n.10, 23–52
- Guillaume D., Dacorogna M., Dave R. Muller V., Olson R., Pictet O. (1997). From the Bird’s Eye to the Microscope: a Survey of New Stylized Facts of the Intra-Daily Foreign Exchange Markets. *Finance and Stochastics* n.1, 95–129.
- Marengo L., Tordjman H. (1996). Speculation, Heterogeneity and Learning: a Simulation Model of Exchange Rate Dynamics. *Kyklos* vol.49, 407–437
- Milgrom P.R., North D.C., Weingast B.R. (1990). The Role of Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the Champagne Fairs. *Economics and Politics* vol.2 n.1, 1–23.
- Rieck C. (1994). Evolutionary Simulation of Asset Pricing Strategies. In E. Hillenbrand and J. Stender (eds.) *Many-Agent Simulations and Artificial Life*. IOS Press, Washington, DC.
- Schleifer A., Summers L.H. (1990). The Noise Trader Approach to Finance. *Journal of Economic Perspectives* vol.4, 19–33.
- Tordjman H. (1998). Some General Questions about Markets. I.I.A.S.A. Interim Report IR-98-025/May.
- Valente M. (1997). Laboratory for Simulation Development: User Manual. I.I.A.S.A. Interim Report IR-97-020/May.