Markov Nets and the NatLab platform; Application to Continuous Double Auction

L. Muchnik[(1)] and S. Solomon[(2)]

1) Department of Physics, Bar Ilan University, Ramat-Gan, Israel
2) Racah Institute of Physics, The Hebrew University, Jerusalem 91904, Israel and Lagrange Interdisciplinary Laboratory for Excellence in Complexity, ISI Turin

ABSTRACT

In describing dynamics of classical bodies one uses systems of differential equations (Newton laws). Increasing the number of interacting bodies requires finer time scales and heavier computations. Thus one often takes a statistical approach (e.g. Statistical Mechanics, Markov Chains, Monte Carlo Simulations) which sacrifices the details of the event-by-event causality. The main assumption is that each event is determined only by events immediately preceding it rather then events in the arbitrary past. Moreover, time is often divided in slices and the various cause and effect events are assumed to take place in accordance to this arbitrary slicing. The dynamics of certain economic systems can be expressed similarly. However, in many economic systems, the dynamics is dominated by specific events and specific reactions of the agents to those events. Thus, to keep the model meaningful, causality and in particular the correct ordering of events has to be preserved rigorously down to the lowest time scale.

We introduce the concept of Markov Nets (MN) which allows one to represent exactly the causal structure of events in natural systems composed of many interacting agents. The Markov Nets preserve the exclusive dependence of an effect event on the event directly causing it but makes no assumption on the time lapse separating them. Moreover, in a Markov Net the possibility exists that an event is affected if another event happens in the meantime between its causation and its expected occurrence. We present a simulation platform (NatLab) which uses the MN formalism to make simulations that preserve exactly the causal timing of events without paying an impossible computational cost.

We demonstrate the use of the MN and NatLab for the study of a continuous double auction market. We ran on NatLab realistic experiments that combined in various market conditions agents with various strategies. The dynamics of the market and the co-evolution of the agents could then be understood on the basis of the detail timing of their actions. Thus, we were able to understand both the traders' intentions and the causes for which their intentions lead to the desired result or not.

INTRODUCTION AND BACKGROUND

What's the problem?

The present paper describes the application of the "Markov Net" (**MN**) concept on a generic continuous time asynchronous platform (**NatLab**) [2] to the study of continuous double-auction markets. The continuous double-auction market is a fundamental problem in economics. Yet most of the results in economics were obtained for markets that are discrete time, synchronous or both. The markets evolve in a reality in which time is continuous and traders act asynchronously by sending independently bid and ask orders. A change in the sequence of arrival to the market of 2 orders may change the actual price by a significant value and eventually lead to completely different subsequent market development (even if their arrival time difference is arbitrarily small). Thus to describe faithfully the market dynamics one has to insure arbitrary precision in the timing of each event (Fig. 1).

The insistence on preserving in great detail the causal order of the various events might seem pedantic and irrelevant for large disordered systems of many independent traders. After all, statistical mechanics is able to reproduce equilibrium thermodynamic results of some Ising-like models without even considering the details of their dynamics [3]. However, there are certain systems in which this requirement seems unavoidable. Economic Markets is one of them.

The classical theoretical framework for understanding market price formation is discrete and synchronous. It is based on the equilibrium price concept introduced by Walras in 1874. Walras' equilibrium price did not address the issue of market time evolution. He assumed that in order to compute the market price it is sufficient to aggregate the offer and demand of **all** the traders. The equilibrium price will then be the one which insures that there is no outstanding demand or offer left.

The extension of the Walrasian equilibrium-price mechanism to time evolving markets was suggested only later by Hicks [4, 5], Grandmont [6] and Radner (1972) [7]. At each time, the demand curves of all the traders were collected and aggregated synchronously. The intuition that ignoring the asynchronous character of the market may miss its very causality, occurred to a number of economists in the past. In fact [21] suggested that economic phenomena like the business cycle may be due exactly to the iterative process of firms / individuals reacting to each other positions. This insight did not have wide impact because the methods used in modeling markets with perfect competition (General Equilibrium theory) and markets with strategic interaction (Game Theory) are synchronous and, arguably, even timeless[1].

If the extension of the Walrasian paradigm[8, 9] to time evolving market price would hold, the sequence in which the orders arrive would be irrelevant and the

[1] We are greatly indebt to Martin Hohnisch for very illuminating and informative correspondence on these points.

insistence of absolute respect of the time order of the various market events unnecessary. However, in spite of the conceptual elegance of the Walrasian construction, its application to time-dependent markets is in stark contrast with what one experiences in the real market. In reality the demand and offer order stacks include at each time only a negligible fraction of the shares and of the share holders.

Moreover, only the orders at the top of the stack have a direct influence on the price changes. One may hope that the local fluctuations due to the particular set of current traders would amount to only some noise on top of a more fundamental dynamics closer to the Walrasian paradigm. However, there are many indications that this is not the case.

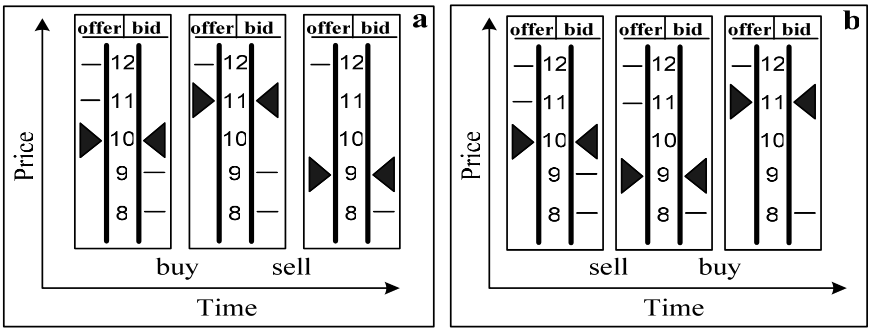

Fig. 1. The figure represents 2 scenarios a) and b) in which 2 traders react by a buy and respectively a sell order to the market having reached the price 10. This moment is represented by the leftmost order book in both a) and b) sequences. The current price is marked by the two arrow heads at 10. We took for definiteness both offer and demand stacks occupied with equal size limit orders (represented by horizontal lines) at integer values around 10. The a) sequence illustrates the case where the seller is faster and saturates the buy order at the top of the buy stack: 9 (this is shown by the order book in the middle). Then, the buyer saturates the lowest offer order: 11 (rightmost column). Thus the final price is 11. The b) figure illustrates the case where the buyer is faster. In this case the final price is 9. Thus an arbitrary small difference in the traders' reaction times leads to a totally different subsequent evolution of the market. The problem is to device a method that insures arbitrary precision in treating the timing of the various events (see discussion of Fig. 13 below).

First, the multi-fractal structure of the market implies that there is no real separation between the dynamics at the shortest time scales and the largest time scales [17, 18]. Thus one cannot indicate a time beyond which the market is bound to revert to the fundamental / equilibrium price.

Second, the motivations, evaluations and decisions of most of the traders are not in terms of an optimal long range investment but rather in terms of exploiting the short term fluctuations whether warranted or not by changes in the "real" or "fundamental" value of the equity.

Third, the most dramatic changes in the price are taking place at time scales of maximum days; so aggregating over longer periods misses their causality.

Forth, the Walrasian auction requires each trader to define a personal demand function, i.e. defining his conditional bid / offer for any eventual price. In practice, not only the traders do not bother and do not have the necessary notions to make their decisions for arbitrary improbable prices, but also the market microstructure just does not have the instruments capable to collect this kind of orders.

All the arguments above can be summed in a more formal tautological form: the variation of the price for large time periods is just the sum (in fact product) of the (relative) variations of the market at the single transaction level. A full understanding of the market dynamics is therefore included in the understanding of the single-trade dynamics. This idea has be investigated in models of high frequency financial data [10, 11, 12] .

The fact that Walras did not emphasize this point is due not only to the lack at the time of detailed single-trade data. Ideologically he had no way to be primed in this direction: the program of Statistical Mechanics of deducing global average thermodynamic properties from the aggregation of the binary interactions of individual molecules was not accepted even in Physics at his time. In recent times there were however trials to reformulate Walrasian ideas in a statistical mechanical framework [13]**.**

The background of the "Markov Net" concept

One may be surprised at the long time that the natural continuous time asynchronous reality have been represented in computers in terms of discrete time synchronous approximations. Even after this became unfashionable, the lack of consideration for the precise causal order of events continued to prevail through the use of random event ordering. Alternatively, in many applications the exact order in which the various events take place is left to incontrollable fortuitous details such as arbitrary choices by the operating system or the relative work load on the computer processors. This is even more surprising given the fact that some fundamental scientific problems fall outside the limits imposed by these approximations. An outstanding example is the ***continuous time double auction market*** that we present in the second part of this paper.

A possible explanation for this state of art is the influence that the Turing discrete linear machine paradigm has on the thinking about computers. This lead to the implicit (thus automatic) misconception that discrete machines can only represent discrete-time evolution. The continuous time is then viewed as a limit of arbitrary small time steps which can be achieved only by pain-staking efforts. We will show that this is not necessarily the case. We also are breaking-away from another implicit traditional assumption in computer simulations: the insistence to see the time evolution of the computer state as somewhat related to the time evolution of the simulated system. Even when the computer time intervals are not considered proportional to the simulated system time, the order in which the events are computed is assumed to be the one of the simulated world. As we shall see, in our implementation of the Markov Nets, there is occasional lack of correspondence between the 2 times. For instance one may compute the time of a later (or even eventually un-happened) event before implementing the event that happens next according to the simulated world causality.

The "Markov Net" **(MN)** idea is easier to understand if one releases oneself from the psychological tendency to assign to the simulated system 1) the discreteness and 2) the causality of the computer process simulating it. The mental effort to overcome this barrier is worthwhile: we can (without computational load penalty) go easily to any time precision (e.g. if desired, one can use time resolution of the order of the time quanta - the time necessary for light to traverse a Plank length ~ 10^{-44} sec in simulations that involve time scales relevant to daily trading activities). The only requirement would be to use double – or if desired arbitrary – real variable bit-length [20]. Using the NatLab framework, one can apply the **MN** concept to applications in heterogeneous agent based macro-economic models [19].

Markov Net Definition and NatLab Description

Definition, Basic Rules and Examples

In this chapter we describe sequences of events in the Markov Net formalism and the way the NatLab platform implements them. While ostensively we just describe a sequence of simulated events, this will allow us to make explicit the way the MN treats the scheduling and re-scheduling of future events as the causal flow is generated by the advancing of the process form one present event to the next.

A Markov Net is a set of events that happen in time to a set of agents. The events cause one another in time with a certain lag between the cause and effect. Thus at the time of causation the effect event is only potential. Its ultimate happening may be affected by other events happening in the meantime. To take this into account, after executing the current event of the Markov Net, the putative happening times of all its directly caused potential effects are computed. The process then jumps to the earliest yet un-happened putative event which becomes thereby the current event. From there, the procedure is iterated as above.

Note that with the present definition, Markov Nets are deterministic structures. We leave the probabilistic Markov Nets study for future research. The name of Markov Nets has been chosen in order to emphasize the fact that the present events are always just the result of their immediate cause events (and not of the entire history of the Net). However, as opposed to a Markov chain [14 , 15], the current state is not dependent only on the state of the system immediately preceding it. For instance, as seen below, the current event could be directly caused by an event that took place arbitrarily far in the past. This is to be contrasted with even the n-order generalizations of the Markov chains where only a limited time strip preceding the present affects it directly. To recover continuum time in an n-order Markov Chain, the time step is taken to 0 and the dynamics collapses into n-order stochastic differential equation (continuous Markov chain) [16]. By contrast, a Markov Net exists already in continuum time and the past time-strip that influences the present is undefined depending on the system dynamics. In this, it shares properties with the Poisson process (except that it allows for a rather intricate causal structure). Thus, only in special regimes it can be approximated by differential equations.

In order to explain the functioning of MN, one represents the various agents time evolution as directed horizontal lines going from some initial time (at the left) towards the future (the right) (Fig. 2).

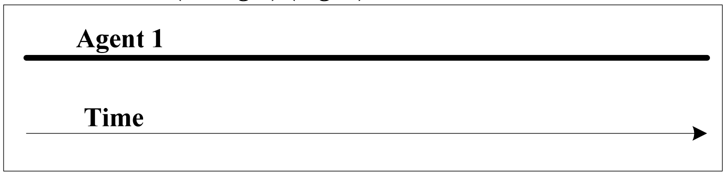

Fig. 2. The line representing an agent time evolution in a Markov Net. Below it, the time axis is plotted too).

The events of receiving some input (being caused) or sending some output (causing) will be represented by points on these horizontal lines. Thus each event is associated with an agent line and with a specific time. Its position on the agent axis will correspond to the Markov Net time of the event (Fig. 3).

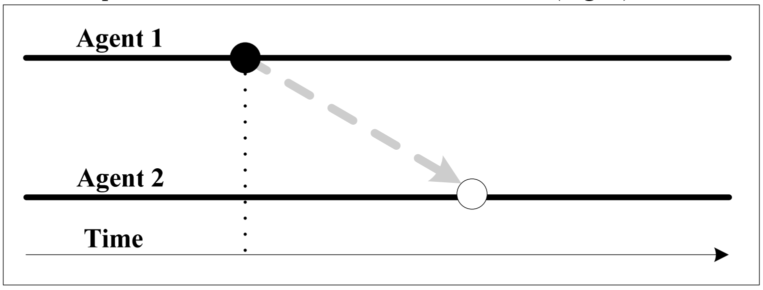

Fig. 3. The black dot on the line of Agent 1 represents the currently happened event. The vertical dotted line represents the time at which the Markov Net evolution currently arrived. The dashed arrow represents the potential causation of a future event (empty dot on the line of Agent2) by the current event.

A vertical line will indicate the current time. Note that in a Markov Net one jumps from the time of one event to the next one without passing through or even considering any intermediate times. The actual transmission / causation will then be represented by a directed line starting on the sending (causing) agent's axis at the sending (cause event) time and ending on the receiving agent axis at a (later) reception (effect event) time (Fig. 3).

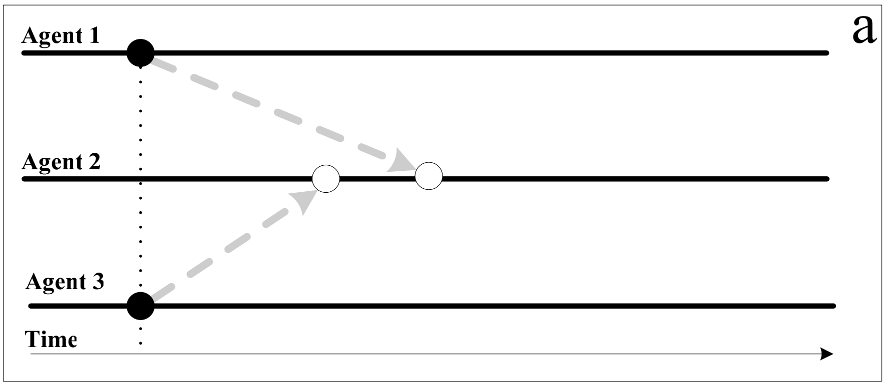

Fig. 4. This Markov Net represents 2 initial events (black dots) belonging to Agents 1 and 3 that cause 2 potential effect events to Agent 2.

Note that while the Markov Net is at a given time (event) the process is (re-) calculating and (re-)scheduling all the future potential events directly caused (or modified) by it. These events may be later modified or invalidated by meantime events so we represent them by open dots pointed to by grey dotted (rather then full black) lines. Those lines are only to represent that the events timings were computed and the events were scheduled, but in fact their timing or even occurrence can be affected by meantime secondary effects of the current state Fig. 5 - Fig. 7). Incidentally, this procedure which is explained here for deterministic events can be extended to probabilistic ones.

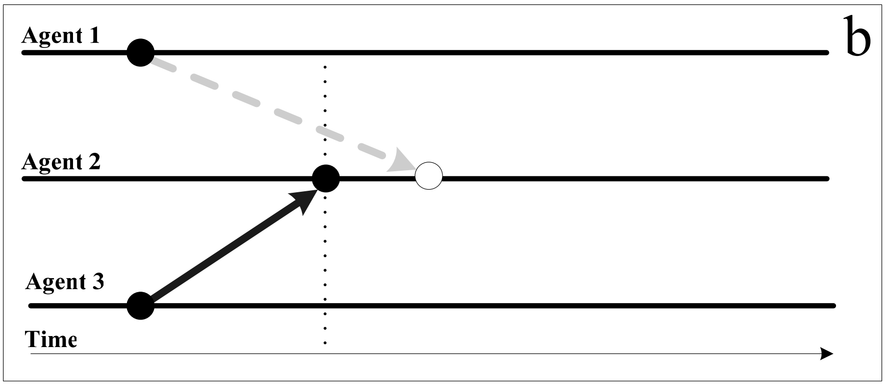

Fig. 5. The Markov Net of Fig 4 advances to the first event on the Agent 2 axis.

The actual time that the transmission / causation takes is computed based on the on the state of the sending (causing) and receiving (affected) agents. Other time delays e.g. between the arrival of an order to the market and its execution are similarly represented.

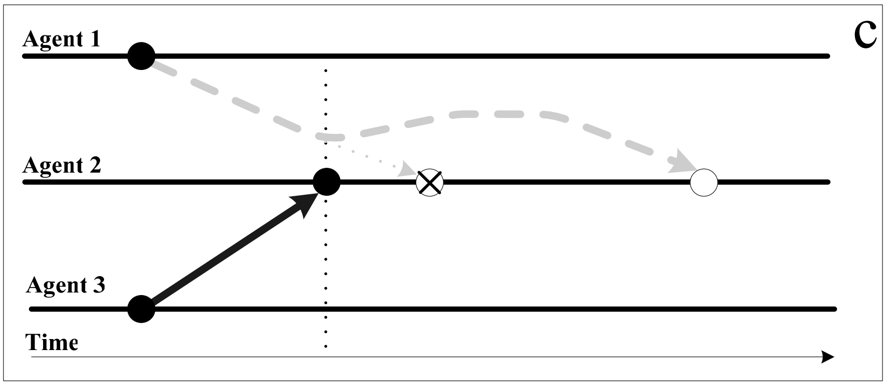

Fig. 6. As a consequence of the first event acting on Agent 2, the second event is modified and remised to a later time.

Since the transmission / causation takes time, it often happens that as a signal travels from one event (agent) to another, another event on the receiver site affects the reception (delaying it, canceling it or changing its effect).

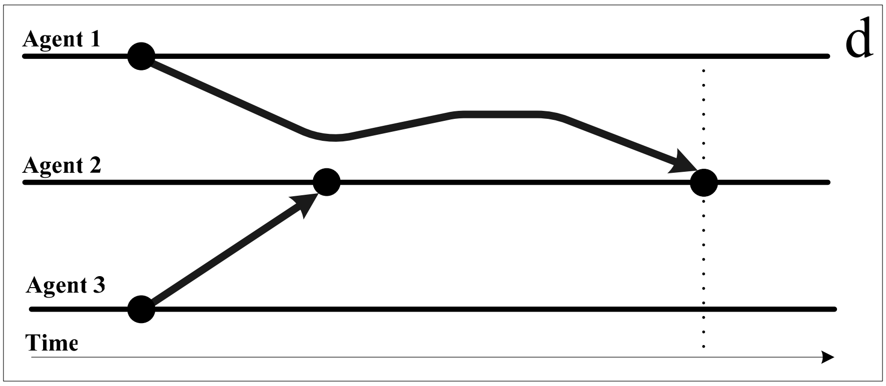

Fig. 7. The modified second event acting on Agent 2 is executed.

The entire dynamics of the system can then be expressed in terms of its Markov Net: the communication / causation lines, their departure (cause) and arrival (effect) event times.

We consider the internal processes of an agent (e.g. thinking on a particular input to produce a decision as output) as a particular case in which an arrow starts and ends on events belonging to the same agent. The beginning of the arrow represents the time at which the process started and the end of the arrow represents the time of conclusion (typically by a decision and sending a communication). The computation of the thinking time and the scheduling of the issue time for the decision is made immediately after the execution of the event that **triggers** the decision process. Of course as with other future events evaluation, if other meantime events affect the agent, the thinking process and the decision can be affected or even cancelled.

Thus the scheduling of an event arrival is always provisory / tentative as the relevant agent can undergo another “unexpected” event before it. Such a intervening event will modify the state of the agent and in particular the timing or even the actual happening of the event(s) that the agent is scheduled to undergo otherwise.

For instance (Fig. 8) , as a result of certain “news” at 15:00, the agent is tentatively computing a certain internal deliberation which ends up with a sell order at 17:00.

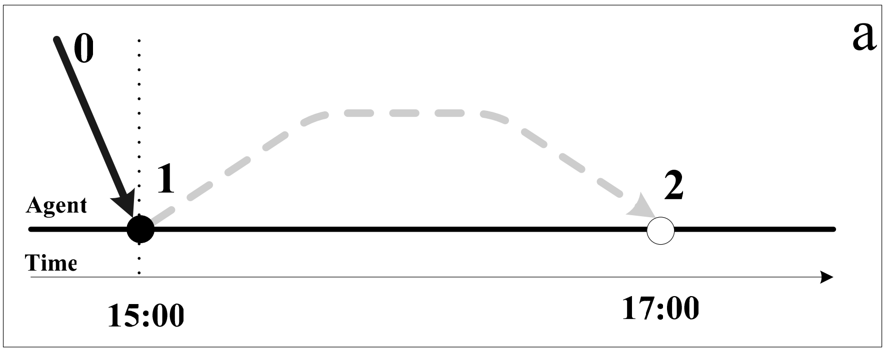

Fig. 8. News arriving at 15:00 cause potentially a decision at 17:00.

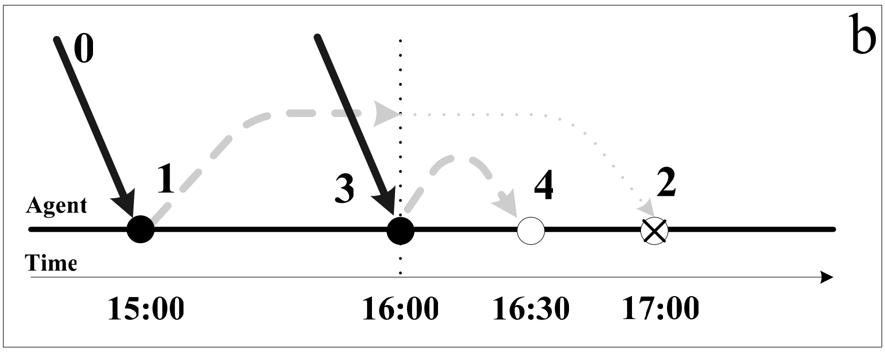

Fig. 9. The news are cancelled by a new “news” at 16:00. The new “news” produce a potential effect at 16:30.

If however at 16:00 the agent receives a message canceling the “news”, the ordering event tentatively scheduled for 17:00 is cancelled (Fig. 9) and some other event (going home at 16:30) can be tentatively scheduled (Fig. 9) and executed (Fig. 10).

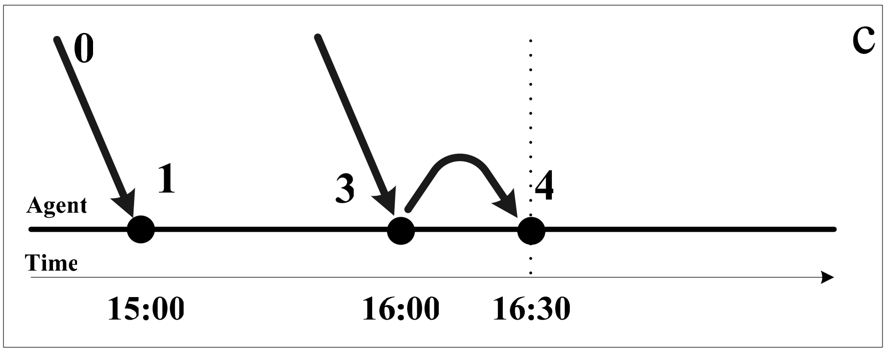

Fig. 10. In the absence of other meantime events, the effect event at 16:30 is actually executed.

Another example of a Markov Net involving 3 agents is shown in Fig. 11. Their interaction sequence is explicitated through the MN scheduling and re-scheduling mechanisms. At the beginning of the process 2 events are pre-scheduled to take place at times t1 (affecting the middle agent) and t2 (affecting the upper agent). These events cause the 2 agents to enter internal deliberation states to prepare some reaction. The middle agent succeeds to make very fast the decision (at t3) and cause the lower agent to send (at t4) a new signal to the upper agent. This new signal, (received at t5) interrupts the internal deliberation state in which the upper agent was. Thus scheduled decision event that was expected to happen at t6 as a result of t2 is cancelled (this is symbolized by the dashed arrow). Instead, as a result of the signal received from the lower agent, the upper agent enters a new internal deliberation that concludes into an answer being sent (at t7) to the lower agent.

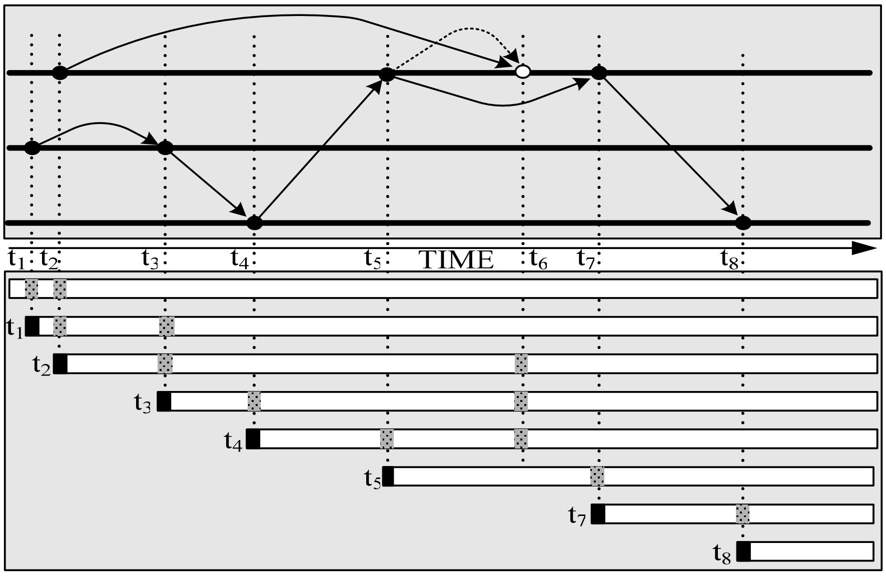

Fig. 11. A Markov Net with 3 interacting agents and interfering events.

Applications to Continuous Double Auction

Simple Example of a Double Auction Markov Net

The simplest MN diagram of a continuous double-auction is shown in Fig. 12. In this diagram we show only one trader (the upper horizontal line) and the market (the lower horizontal line). The market price evolution is shown below the market line. The process described by this diagram starts with the agent setting P0 (represented on the price graph by a *) as an attention trigger. This is represented by a formal (0-time delay) message sent to the market (at time t1). When the price reaches P_0 (at time t2), the agent notices it (t3), observes it (t4) and makes the decision to send a buy order (t5) which will be executed at t6. At the same time (t5), the agent also sets a new attention trigger price P_1.

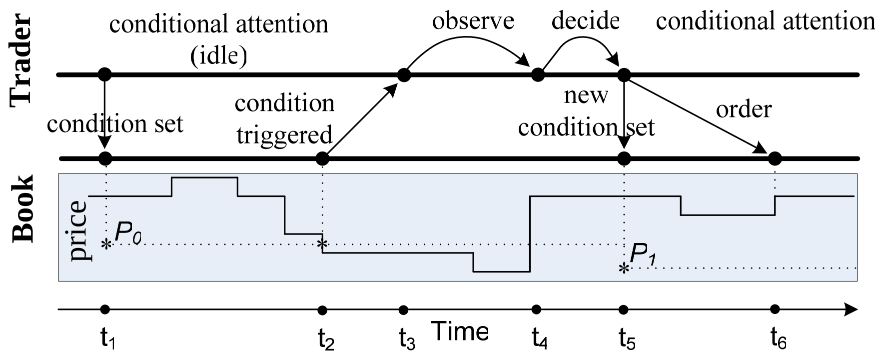

Fig. 12. Schematic implementation of the continues double auction

Returning to the example in the introduction, we can now see how the Markov Net formalism takes care of the details of the traders' behavior [25, 26, 27] in order to represent faithfully the sequence of market events that their behavior engenders.

In Fig. 12 one can see that at certain past times (t1 and t2), each of the traders have fixed 10 as an action trigger price threshold. This is represented in the MN conventions by messages that each of them sent to the market. This is only an internal NatLab platform technicality such that the messages are considered as not taking any transmission time in the simulated world. As a consequence of these triggers, when the market reaches 10 at time t3 (following some action by another trader), each of them receives a "message" which is the representation on the platform of the fact that their attention is being awakened by the view (presumably on their computer screen) of the new price display equal to 10.

Due to the delays in their perception (and other factors), their reception times (t4 and t5) are slightly different. Once they perceive the new situation they proceed to think about it and reach a decision. They have different times until they reach the decision. Thus in the end they send their market orders at different times (t6 and t7). After taking into account the time each order takes to reach the market, one is in the position to determine the arrival times of each of the orders (t8 and t9). Thus, for given parameters (perception times, decision times, and transmission lags) those times are reproducibly computable and unambiguous market scenarios can be run. Fig. 13 A and B illustrate the 2 possible outcomes discussed in Fig. 1. Incidentally, upon seeing the effect of the order of Agent 2, Agent 1 may wish to take fast action to cancel his own order. Of course this possibility depends on the existence of a second communication channel to the market that is faster then the one on which the initial order has been sent.

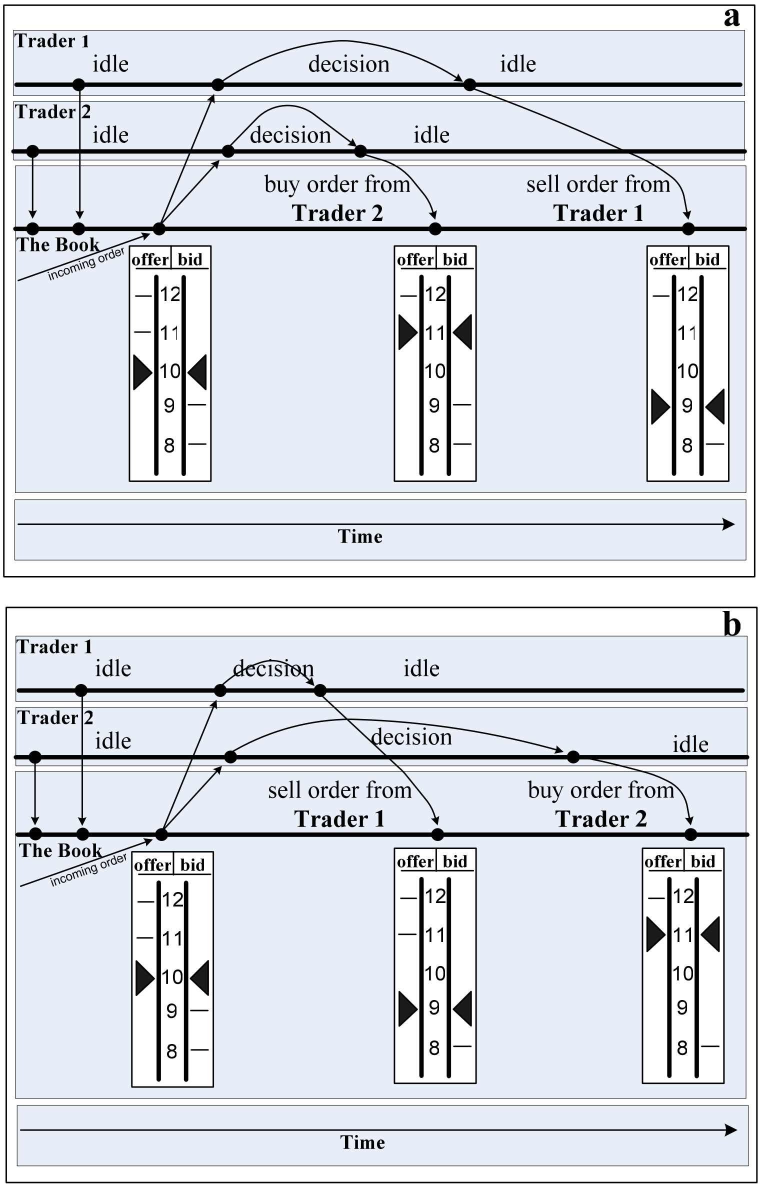

Fig. 13. The order books are the same as in Fig. 1. The figure describes in detail the Markov Net implementation of the correct time order in such a situation. The upper (lower) panel describes the situation in which the buy (sell) order arrives first corresponding to the a) (b) sequence in Fig. 1.

Application of Markov Nets to Realistic Continuous Double-Auction Market

We have performed computer-based experiments of a continuous double-action market. As seen below (e.g. the spectacular loss of the professional strategy 7) the details of the timing of the various events at the lowest time-sacel (single trade) have a overwhelming influence on the market outcome.

The NatLab experiments were based on the “Avatars” method that we have introduced lately [ref Gilles] in order to insure that our simulations / experiments are realistic. The “Avatars” method allows us to capture in the computer the strategies of real traders (including professionals). By doing so we are crossing the line between computer simulations and experimental economics [1]. In fact we have found that NatLab is a convenient medium to elicit, record, analyze and validate economic behavior. Not only NatLab provides in a most motivating and realistic environment but also provides a co-evolution arena for individual strategies dynamics and market emergent behavior.

```
//--------------------------------------------------------------------------
TOrder TTrader15::IssueOrder()
{
    TOrder NewOrder;
    NewOrder.ID          =   ID;
    NewOrder.IssueTime   =   Manager.CurrentVirtualTime;
    NewOrder.OrderType   =   otNoOrder;
    NewOrder.Price       =   Manager.CurrentPrice;
    NewOrder.Shares      =   TSharesType(0);

    double AvgK =     ( (Manager.History->Size) ?
        Manager.History->AveragePrice(Manager.CurrentVirtualTime-TVirtualTime(1000)) :  Manager.CurrentPrice.RealPrice);
    double AvgD =     ( (Manager.History->Size) ?
        Manager.History->AveragePrice(Manager.CurrentVirtualTime-TVirtualTime(100)) :  Manager.CurrentPrice.RealPrice);
    double StdK =     ( (Manager.History->Size) ?
        Manager.History->StdvPrice(Manager.CurrentVirtualTime-TVirtualTime(1000)) :  Manager.CurrentPrice.RealPrice);

    if (AvgD > AvgK + StdK) { // SELL! The short-term average is below the long-term one. the price went up.
           NewOrder.Price.TickNumber = (NewOrder.Price.TickNumber > 1) ? NewOrder.Price.TickNumber-1 : NewOrder.Price.TickNumber;
           NewOrder.Shares      =   TSharesType(0.2* Shares.Shares);     // sell 20% of the shares
           NewOrder.OrderType   =   otSellLimit;
    }
    else if (AvgD < AvgK - StdK) { // BUY! The short-term average is below the long-term one. the price went down.
            NewOrder.Price.TickNumber = NewOrder.Price.TickNumber + 1;
            NewOrder.Shares      =   TSharesType(0.1* Cash.RealPrice/NewOrder.Price.RealPrice);   // buy with 10% of available cash
            NewOrder.OrderType   =   otBuyLimit;
    }
    if (NewOrder.Shares.Shares==0) { NewOrder.OrderType = otNoOrder; }
    return NewOrder;
}
//--------------------------------------------------------------------------
```

Fig. 14. The NatLab code for the long range mean reverters (trader group number 4 in the text).

Casting the subjects’ strategies within computer algorithms (Avatars) allows us to use them subsequently in various conditions: with and without dividend distribution, with and without interest on money in the presence or absence of destabilizing strategies. We present below in detail some experiments.

We have collected the preferred trading strategies from 7 subjects and we have created (by varying the individual strategies parameters) a set of 7000 trading agents. We have verified that the subjects do agree that the traders behave according to their intentions. We then made runs long enough to have a number of trades per trader of the order 1000. In addition we introduced 10000 random traders to provide liquidity and stability. One can also consider their influence on the market as a surrogate for market makers (which we intend to introduce in future experiments).

In this particular experiment we verified the influence on the market of having 1000 of the traders respecting a daily cycle. We did so by first running the market simulation in the absence of the periodic traders and then repeating the runs in their presence.

Let us first describe some of the relevant traders, their algorithms and their reaction to the introduction of the periodic daily trends (the numbers on the curves in the graphs in Fig. 15 to Fig. 18 correspond to the numbers in the following list).

1. Random

We introduced random traders. Each random trader extracts his belief about the future price from a flat, relatively narrow distribution around the present price and then offers or bids accordingly with a limit order set to be executed at the price just below (offer) or above (bid) the current market price,. The order volume is proportional to the number of agent's shares (in case of offer) or to his cash (in case of bid) and to the relative distance between the current market price and the price the agent believes is right.

As seen on Fig. 15 the random traders are performing poorly; their average (Fig. 16) is very smooth because their guesses at each moment are uncorrelated. Their performance and behavior are not affected by the presence or absence of daily trends due to their inability to adjust to them (Fig. 17, Fig. 18).

2. Short range mean reverters

They compute the average over a certain - very short - previous time period and they assume that the price will return to it in the near future. On this basis they buy or sell a large fraction of their shares.

When executed on the market with no periodic daily trends (Fig. 15), this strategy took full advantage of the observed short-range negative autocorrelation of price returns. In fact, it was the winning strategy in that case. Agents following this strategy had their behavior strongly correlated among themselves. They were able to adjust to the continueously changing fluctuations (Fig. 16).

When periodic trends (and hence short-range correlations in returns annihilating the natural anti-correlations) were introduced (Fig. 17), the performance of the strategy dropped dramatically. However, even in this case, positive returns were recorded. In the case of the periodic market, one can clearly identify the emergence of correlation between the actions of all agents in this group (Fig. 18).

3. Evolutionary extrapolation

Each agent which follows this strategy tries to continuously evaluate its performance and occasionally switches between trend following and mean reverting behavior.

Due to relatively long memory horizon any trend and correlation that could have been exploited otherwise was averaged out and the strategy shows moderate performance (Fig. 15, Fig. 17) with almost no spikes of consistent behavior (Fig. 16, Fig. 18).

4. Long range mean reverters

They compare long range average to short term average in order to identify large fluctuations and submit moderate limit orders to exploit them.

This strategy does not have any advantage in the case when no well-defined fluctuations exist and can not outperform other strategies (Fig. 15). The actions of the 1000 agents following it are uncorrelated. and their performance is uncorrelated and in general is averaged to 0 (Fig. 16).

However, when long trends do appear in the case when periodic fluctuation of the price are induced, agents following this strategy discover and exploit them (though, with some delay) (Fig. 18) to their benefit. In fact, they expect that the price fluctuation will end and it will eventually revert back to the average. Unlike the short-range mean reverters they are not disturbed by the small fluctuations around the - otherwise consistent – longer cycle. Only small investment volumes (due to their high risk avertion) do not allow the agents using this strategy to outperform the best strategy in the daily periodic case (Fig. 17).

5. Rebalancers+Trend Followers

They keep 70% of their wealth in the riskless asset. The rest of 30% is used in large order volumes in a trend following manner: if the short-range average price exceeds the long range average by more then one standard deviation they buy. If the price falls by more then one standard deviation, they sell. Agents practicing this strategy use market orders to ensure that their intentions are executed.

In general this strategy is relatively immune to sudden price fluctuations since it causes its agents to put at risk at most 30% of their wealth. It attempts to ride moderate trends. This does not produce remarkable performance in the first case when trends are nor regular. The strategy performance is average in this case.

However, it is the best possible strategy in the case of periodic fluctuations. The agents following this strategy, discover the trends quickly and expect them to continue for as long as the price evolution is consistent with them. In fact, they consistently manage to predict the price evolution and make big earnings (Fig. 17).

6. Conservative mean reverter

Similar to 4) but with slightly shorter range average and less funds in the market. This does not make any difference when no fluctuations are present (Fig. 15). However, when the periodic fluctuations are introduced (Fig. 17) they expect the price to return back to the average sooner that it actually does. Thus they are mistaken for rather long intervals.

7. Professional volume watcher

They watch the aggregated volumes of the buy and sell limit orders. If a large difference towards offer is discovered, they sell and viceversa. Their intention is to predict starting dynamics and respond immediately upon smallest possible indicators by issuing market orders so that the deal is executed ASAP.

This strategy though suggested by a professional trader (and making theoretically some sense) fails to perform well in both cases. The reason for that is, in our opinion, the wrong interpretation of the misbalance in the orders. Agents of this type expect the volume of the limit orders in the book to indicate the excess demand or supply and hence give a clue on the immediate price change. However, in our case, the volumes of limit orders on the both sides of the book are highly influenced by big market orders (for example, those issued by 5) and in fact are the result of very recent past large deals rather then a sign for a incoming trend. Another way to express this effect is to recall that the prices are strongly anticorrelated at one tick. Thus this strategy is consistently wrong in it's expectations. Interestingly enough, its disastrous outcome would be hidden by any blurring of the microscopic causal order (that would miss the one-tick anticorrelation).

8. The Daily traders:

They gradually sell all their holdings in the “morning” and then re-buy them back in the “evening hours”. This causes the market price to fluctuate accordingly.

What is more important, the behavior of all the other strategies is greatly affected as well as their relative success. Some of these strategies adapt and exploit the periodic trends while others are loosing systematically from those fluctuations.

Some of the results (specifically the time evolution of the average relative wealth of the traders representing each subject) are plotted in Fig 15 and Fig 17. The unit of time corresponds to the time necessary to have in average a deal per each agent.

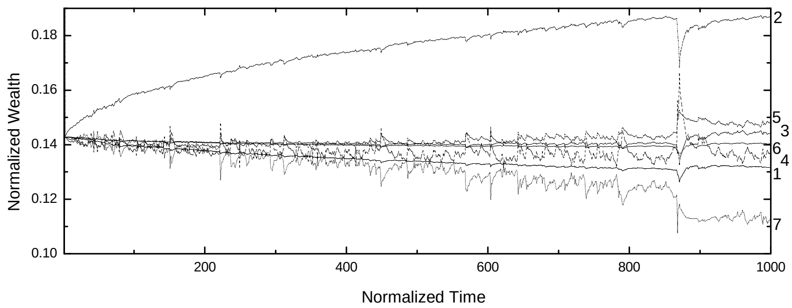

Fig. 15. Relative wealth in a market including the first 7 types of traders. The strategy labels on the left border of the graph correspond to the numbers labeling the traders in the text.

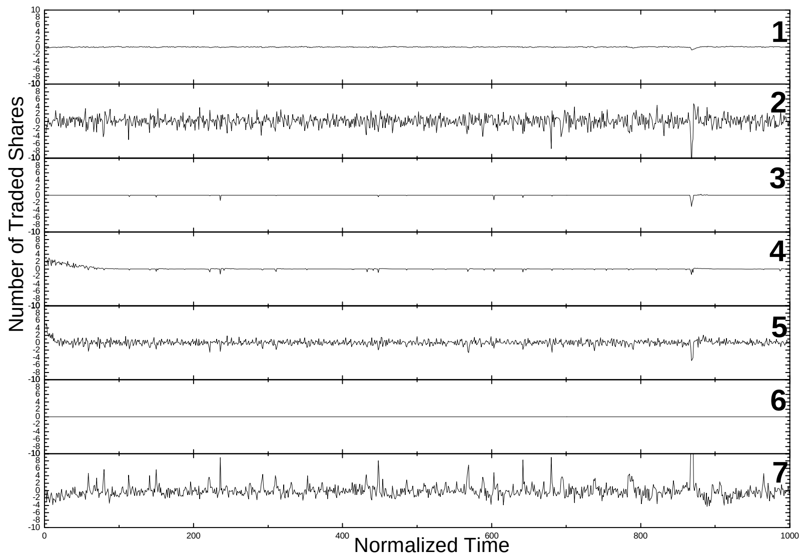

Fig. 16. The shares trading history is plotted for each strategy. The average number of bought shares per time unit per trader is represented by the height of the graphs above 0. The number of sold shares is represented by the depth below 0.

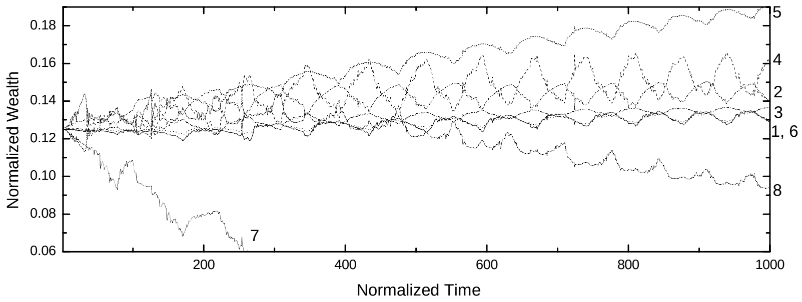

Fig. 17 The relative wealth evolution in a market containing all types of traders including daily traders (number 8).

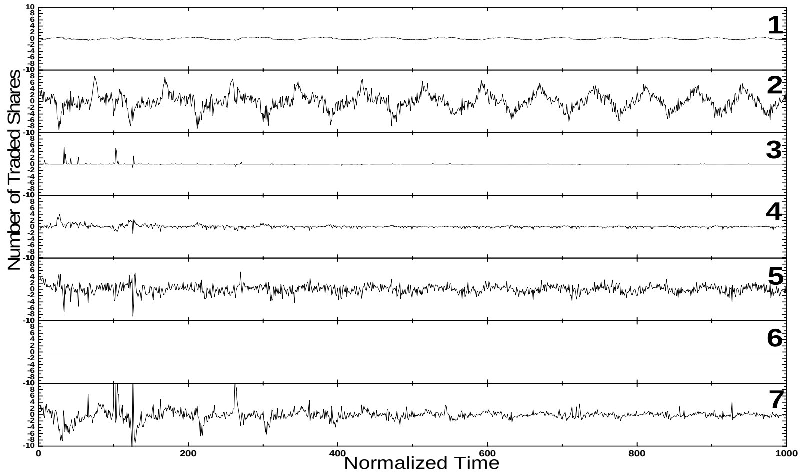

Fig. 18. The shares trade as in the Fig 16, but for the run with daily trends.

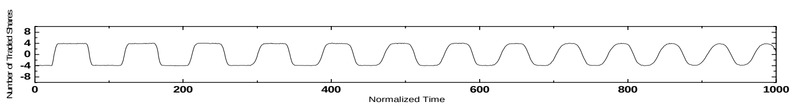

Fig. 19. The activity of the daily trader (strategy number 8 in the list)

Conclusions and Outlook

For a long while, the discrete time Turing machine concept and the tendency to see computers as digital emulations of the continuous reality lead to simulation algorithms that mis-represented the causal evolution of systems at the single event time scale. Thus the events took place only at certain fixed or random times. The decision of how to act at the current time was usually by picking up (systematically or randomly) agents and let them act based on their current view the system (in some simulations, event ordering was even left to the arbitrary decision of the operating system!). This neglected the lags between cause, decision, action and effect. In the "Markov Net" representation it is possible for an event to be affected by other arbitrary events that were caused in the meantime between its causation and its happening. This is achieved exactly and without having to pay the usual price of taking a finer simulation mesh.

We have constructed a platform (NatLab) based on the Markov Net (MN) principle and performed a series of numerical and real-subject experiments in behavioral economics. In the present paper we have experimented with the interactions and emergent behavior of real subjects' preferred strategies in a continuous double-auction market.

In the future, we propose to extend the use the NatLab platform in a few additional directions:

- Experiment with the effect on the market of various features and events.
- Compare the efficiency of different trading strategies.
- Isolate the influence of (groups of) traders' strategies on the market.
- Study the co-evolution of traders' behavior.
- Find ways to improve market efficiency and stability.
- Study how people depart from rationality
- Study how out-of-equilibrium markets achieve or not efficiency
- Anticipate and respond to extreme events
- Treat spatially extended and multi-items markets (e.g. real estate), firms dynamics, economic development, novelty emergence and propagation.
- Study the effect of out of market communication, and coupled networks [22]

Last but not least, the mathematical properties of Markov Nets are begging to be analyzed.

Acknowledgment

We are very grateful to Martin Hohnisch for sharing with us his extensive and deep understanding of the economic literature on price formation mechanisms. We thank David Bree, Gilles Daniel, Diana Mangalagiu, Andrzej Nowak, Enrico Scalas, for very useful comments. We thank the users of the NatLab platform for their enthusiastic cooperation. This research is supported in part by a grant from the Israeli Academy of Sciences. The research of L.M is supported by Yeshaya Horowitz Association through The Center for Complexity Science. The present research should be viewed as part of a wider interdisciplinary effort in complexity that includes economists, physicists, biologists, computer scientists, psychologists, social scientists and even humanistic fields [22, 23, 24]

References

1. Daniel G, Muchnik L, and Solomon S (2005) "Traders imprint themselves by adaptively updating their own avatar" Proceedings of the 1st Symposium on Artificial Economics, Lille, France, Sept. 2005
2. Muchnik L, Louzoun Y, Solomon S (2005) Agent Based Simulation Design Principles –Applications to Stock Market, "Practical Fruits of Econophysics", Springer, Verlag Tokyo, Inc
3. McCoy B M, Wu T T (1973), The Two-Dimensional Ising Model, Harvard University Press (June 1, 1973), ISBN: 0674914406
4. Hicks J R, (1939), Value and Capital, (Oxford, England: Oxford University Press).
5. Hicks J R, (1946), Value and capital : an inquiry into some fundamental principles of economic theory, Oxford: Clarendon, 1946.
6. Grandmont J-M, (1977) "Temporary General Equilibrium Theory," Econometrica, Econometric Society, vol. 45(3), pages 535-72.
7. Radner, R. (1972): .Existence of Equilibrium of Plans, Prices and Price Expectations in a Sequence of Markets,. Econometrica, 40, 289.303
8. Walras L, Elements of Pure Economics (London: George Allen and Unwin, 1954 [1874]).
9. Walras L, (1909) "Économique et Méchanique", Bulletin de la Société Vaudoise 45, p.313-325
10. Scalas E, (2005) "Five Years of Continuous-time Random Walks in Econophysics," Proceedings of WEHIA 2004, to be published. http://econwpa.wustl.edu/eps/fin/papers/0501/0501005.pdf
11. Scalas E, Gorenflo R, Luckock H, Mainardi F, Mantelli M, Raberto M, (2004), "Anomalous waiting times in high-frequency financial data", Quantitative Finance 4, 695-702, 2004. http://tw.arxiv.org/PS_cache/physics/pdf/0505/0505210.pdf
12. Scalas E, Gorenflo R, Mainardi F, (2000) "Fractional calculus and continuous-time finance ", Physica-A (Netherlands), vol.284, p.376-384, 2000, http://xxx.lanl.gov/abs/cond-mat/0001120
13. Smith E, Foley D K, (2004), Classical thermodynamics and economic general equilibrium theory, http://cepa.newschool.edu/~foleyd/econthermo.pdf
14. Bharucha-Reid A T (1960), Elements of the Theory of Markov Processes and Their Applications. New York: McGraw-Hill, 1960. ISBN: 0486695395
15. Papoulis, A. "Brownian Movement and Markoff Processes." Ch. 15 in Probability, Random Variables, and Stochastic Processes, 2nd ed. New York: McGraw-Hill, pp. 515-553, 1984.
16. Wiener N, "Differential space", J. Math. Phys., 2(1923), 131.

17. Mandelbrot B, (1963). The variation of certain speculative prices. The Journal of Business, 36, 394-419.
18. Liu Y, Gopikrishnan P, Cizeau P, Meyer M, Peng C K, Stanley H E (1999), Phys. Rev. E 60, 1390 (1999).
19. Gallegati M, Giulioni G, (2003) "Complex Dynamics and Financial Fragility in an Agent Based Model," Computing in Economics and Finance 2003 86, Society for Computational Economics.
20. Koren I, (2002) Computer Arithmetic Algorithms, A. K. Peters, Natick, MA, 2002,ISBN 1-56881-160-8
21. Lindahl E, "The Place of Capital in the Theory of Price", 1929, Ekonomisk Tidskrift.
22. http://shum.huji.ac.il/~sorin/ccs/giacs-annex12.doc
23. http://shum.huji.ac.il/~sorin/ccs/co3_050413.pdf
24. http://europa.eu.int/comm/research/fp6/nest/pdf/nest_pathfinder_projects_en.pdf
25. Tversky, Amos, and Daniel Kahneman, ‘‘Judgment under Uncertainty: Heuristics and Biases,’’ Science, CLXXXV (September 1974), 1124–1131.
26. Kahneman, Daniel, Jack L. Knetch, and Richard H. Thaler, ‘‘Fairness and the Assumptions of Economics,’’ Journal of Business, LIX (1986), S285–300.
27. Akerlof, George A., An Economic Theorist’s Book of Tales (Cambridge, UK: Cambridge University Press, 1984).
28. Lindahl E R, (1929),"The Place of Capital in the Theory of Price", 1929, Ekonomisk Tidskrift
29. Erez T, Hohnisch M, Solomon S, (2005), "Statistical Economics on Multi-Variable Layered Networks", in "Economics: Complex Windows", Eds. M. Salzano and A. Kirman, (Springer, 2005), p. 201