AN AGENT-BASED MODEL FOR CRISIS SIMULATION IN PAYMENT SYSTEMS

L. ARCIERO, Bank of Italy
C. BIANCOTTI*, Bank of Italy
L. D'AURIZIO, Bank of Italy
C. IMPENNA, Bank of Italy

ABSTRACT

This paper presents an agent-based model of a Real Time Gross Settlement (RTGS) payment system. Banks are represented as agents who exchange payment requests, and decide on the timing and the mode of settlement based on a set of simple rules. While highly stylized, the model features the main elements of a real-life system, including a central bank acting as lender of last resort, and a simplified money market. Simulations are run to predict the impact of a disruptive event on the flow of interbank payments. When one of the banks participating in the system is hit by such an event, resulting in the impossibility to perform transactions, three distinct phases emerge. The first one is characterized by inflated liquidity expectations, the second one features a thickening of the money market and of payment queues, and the third one is marked by an increase in defaulted obligations. In order to staunch the flow of losses and restore the orderly functioning of the payment system, central bank intervention must not only be timely, but also relatively intense in terms of the amount of liquidity funneled to the system.

Keywords: Agent-based modeling, Payment systems, RTGS, Liquidity, Crisis simulation

INTRODUCTION

In modern exchange economies, the reliability and the efficiency of payment systems represent fundamental pre-conditions for smooth and safe financial transactions of banks, firms, and households.

The value of payments has increased dramatically in the last decade, as a result of financial liberalization, innovation processes, and increasing globalization of the real economy. In the European Union, interbank payments amounted to 57 times annual GDP in 2005, up from 40 times at the end of the Nineties. Given these developments, in the last decade central banks and market participants have been devoting specific attention to the payments settlement phase, where financial risks are more likely to produce potential systemic impacts. The need to manage and mitigate such risks, at the same time facilitating the handling of an increased volume of transactions, has led to the widespread adoption of Real Time Gross Settlement (RTGS) systems, where individual transactions are settled in real time and with immediate finality.

* Corresponding author address: Claudia Biancotti, Bank of Italy, Department of Economic and Financial Statistics, Via Nazionale 91, 00184 Rome, Italy; e-mail: claudia.biancotti@bancaditalia.it.

Such systems require each participant to hold adequate liquidity levels on an intraday basis; should this not be the case, streams of payment operations might go unfulfilled, triggering undesirable domino effects. A disruptive event, be it physical, technical or financial, may induce prolonged illiquidity conditions, with potentially severe consequences for economic activity. It is therefore very important, especially for central banks as promoters of financial stability, to gain an understanding of how these illiquidity conditions arise, which parts of the system they affect most, and which strategies are most effective when attempting to counteract them.

This paper presents an agent-based model aimed at discerning how a traumatic event affecting a single bank at a given time impacts on liquidity levels and expectations of all other banks participating in the same payment system for the rest of the operational day. While by no means exhaustive in representing the complexity of actual RTGS systems, the model incorporates the core behavioral rules of banks under ordinary conditions; emerging post-crisis behavior appears to be consistent with observations of real-life episodes, if still very simplified.

The paper is structured as follows. Section 2 presents a brief overview of the literature on the simulation of payment systems. Section 3 describes the operating principles of a stylized RTGS system. Section 4 provides the detail of our agent-based model. Section 5 discusses the results. Section 6 concludes and puts forth proposals for future work.

SIMULATION METHODS FOR PAYMENT SYSTEMS

The simulation approach is very suitable for the representation of payment systems: it enables researchers to build models closely replicating the real operational environment, unconstrained by the existence of numerous complex interrelations which are typically hard to represent through traditional econometric tools. Simulations provide information both on the normal functioning of the system and on extreme, not frequently observed scenarios. Input data can be of different kinds: time series of payments submitted by banks can be used for "what if" analyses under different settlement mechanisms, whereas stochastic inputs can be either used for theoretical studies or for models aimed at extrapolating the consequences of particular behavioral assumptions on small-scale settings.

These techniques are now a reliable support in designing payment systems that can control their typical risks, as described by the Bank of England (2000): credit, liquidity, operational, and legal. The Bank of Finland pioneered the construction of simulation models by building an *ad hoc* algorithm (Leinonen, 2005): a deterministic stream of payments is accepted as input and dealt with according to different sets of rules. Bank behavior is taken as given, or is made able to evolve in a predetermined manner.[1]

When considering the largish menu of possible simulation methods currently available, the agent-based framework (Gilbert and Terna, 2000; Fioretti 2004) appears to the best option for our task. Payment systems are coherent in a recognizable way, but their elements, interactions, and dynamics generate structures admitting surprise and novelty which cannot be defined *a priori*. They are more of the sum of their parts; also, the time dimension is explicitly relevant in their functioning, in that different aggregate scenarios may emerge according to the

[1] Payment systems can also be represented as a complex network, with banks as nodes and mutual liabilities/claims defining the arches. This kind of modeling is, however, static: in the sense that the time dimension is not directly taken into account. Network theory has been exploited to study the main features of real interrelations among banks (Boss et al., 2004), making it possible to understand the concentration level of the system, i.e. whether few banks are responsible for the bulk of the links. The consequences of catastrophic events are simply modeled by removing a specific node and measuring the performances of the rest of the network.

payment sequence and rules governing interactions. In other words, they fit the definition of complex systems as provided by Wolfram (1994).

In a way that lends itself well to agent-based modeling, the behavior of commercial banks, at least in the short run, can also be represented in terms of a simple and consistent set of rules, governing a core set of decisions. Given a flow of payments, banks mainly choose in which order these payments have to be submitted in the system, and how to obtain the liquidity necessary to meet obligations, subject to known constraints (Markose et al., 2006). The set of available strategies can be described in the language of game theory (Bech and Garrat, 2003), and translated into algorithms with ease. Learning mechanisms can also be implemented, allowing banks to move from a set of rules to another according to the values of an objective function dynamically updated by simulation results. Predictive learning modules can double as tests of whether banks' adaptive behaviors converge to a steady state in terms, for example, of the liquidity committed in the system (Galbiati and Soramäki, 2007).

SOME RELEVANT FEATURES OF RTGS SYSTEMS

Figure 1 describes the basic functioning of an RTGS environment. In the following, we give a sketch only of those features that are especially relevant for our simulation, and do not appear to have been fully considered in any other similar exercise. For a comprehensive illustration of the underlying system, see Arnold et al. (2006).

Independent of the underlying instrument (bills, checks, electronic transfers, etc.), each payment operation in an RGTS context generates an integrated process, going from the initial decision to transfer funds to a counterparty and until the final settlement in central bank money. In such processes, the main role is played by commercial banks: at some stage, payments between customers of different banks are likely to be treated as interbank flows of liquidity

Several intraday liquidity sources are available to banks. Within the system itself, each bank does normally rely on a continuous flow of incoming payments from its counterparties. Moreover, it can obtain central bank intraday credit, which entails a cost either explicit, if such credit is subject to a fee, or implicit, whenever the provision of funds is not priced but is conditional on the availability of collateral. Alternatively, funds can be borrowed from other banks, in the interbank market.

Since liquidity is costly, banks are however involved in a strategic game, possibly affecting the time pattern chosen to send payments. More specifically, they face, on a continuous time basis, a trade-off between liquidity and delay costs. By releasing payments as soon as funds are available they satisfy customer and counterparty needs and benefit from a sound reputation, but can incur high liquidity costs, to the extent that they borrow from the money market or the central bank. On the other hand, banks can more effectively play on the intraday dynamics of the money market by choosing to delay payments, at the expense of increased sytemic uncertainty and worsening of reputation.

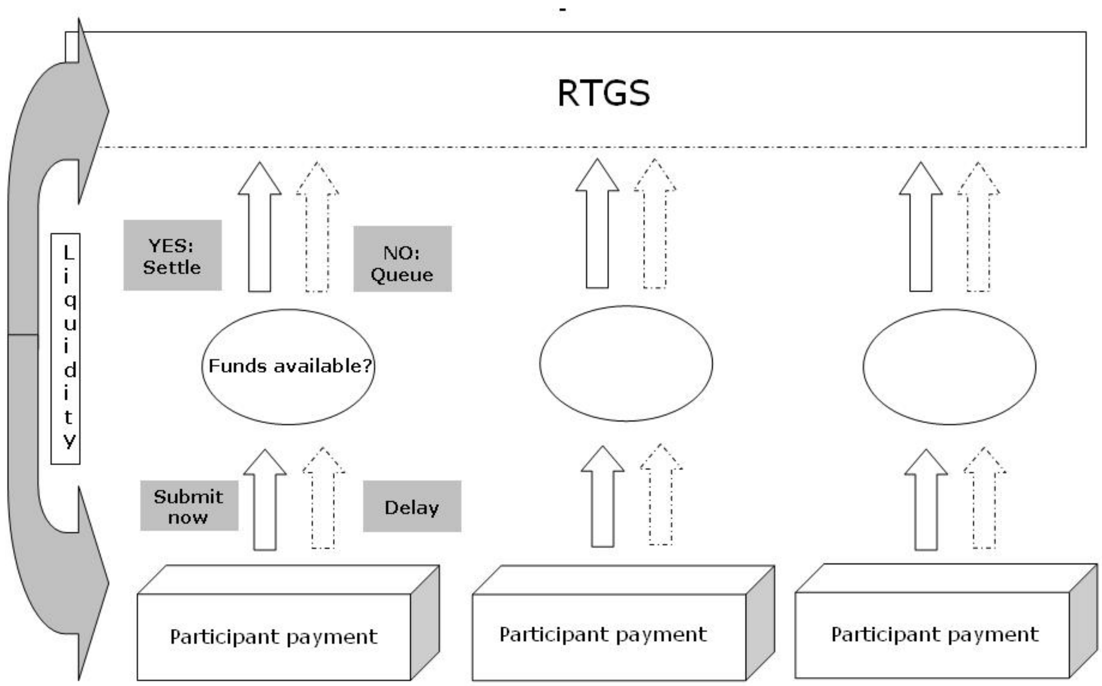

FIGURE 1 Basic functioning of an RTGS environment

DESCRIPTION OF THE MODEL

We model a stylized version of a "plain vanilla" RTGS system, excluding advanced liquidity management tools such as optimization and centralized queues. The model is implemented in the StarLogo[2] environment, and it includes seven breeds of agents: banks, the central bank, payment requests, defaulted operations, interbank loans, crisis events and craters, representing banks hit by such events and accordingly unable to perform any operation. One StarLogo second corresponds to one real-life minute.

During the setup phase, representing the start of the operational day, banks are endowed with a starting level of cash and collateral. During the day, for every tick of the clock, each bank hatches a certain number of new agents, representing aggregations of all payment requests to be delivered to a single counterparty in that moment. For computational reasons, the RTGS system is treated as direct-debit based, with payments always requested by the payee. Payment requests are assigned individual deadlines, ranging from "upon reception" (time-critical payments) to a certain number of minutes after reception. Amounts, deadlines and counterparties are determined through a random draw, whose features can vary depending on the desired scenario.

Each payment request proceeds to cross the StarLogo terrain at fixed speed towards its destination. Upon arrival, it is queued until its deadline expires, and triggers the settlement routine. It is noteworthy to stress that this submission rule allows for the implicit modeling of settlement delay costs banks incur: for each operation, the delay cost is assumed to be a

---

2 Starlogo TNG Preview 4.2, released in April 2007, freely available at http://education.mit.edu/starlogo-tng/.

discontinuous function of time with a jump at the payment deadline, being lower than liquidity costs until the deadline and greater afterward.[3]

From the moment a payment request is generated and until it is settled, the amount thereof is incorporated in the expectations of a liquidity change for both the originating and receiving bank. We therefore assume that in each moment banks are perfectly informed on all payment requests concerning them either as payer or payee; at time t banks are able to calculate their future liquidity up to the moment T, where $T-t$ is the maximum lifespan of a payment request generated at t. The expectation is constantly updated as new payment requests enter the world.

When the settlement process starts, the intended payer tries to meet its obligation with cash. If the cash balance held at the moment is not sufficient, it tries to pledge collateral at the central bank, who provides liquidity based on 100% percentage. If collateral is also insufficient the bank tries to borrow on the money market ("short" bank), thus looking for a lending counterparty ("long" bank). The short bank i randomly draws a potential counterparty k among all the other banks with the exception of j, the bank who originated the payment request p^{ij}_t. Bank k agrees to the loan if the condition $p^{ij}_t \leq E\,[L^k_T]$ is met; $E\,[L^k_T]$ is the expected cash balance for k at time T. In other words, the loan to cover p^{ij}_t is extended on the basis of both present cash balances and future liquidity expectations, as determined by other payments currently in existence initiated by or sent to the potential lender.

Whenever a loan request is refused the payment request bounces, and the short bank looks for another lender. A counter keeps track of the number of bounces per request: when they exceed a certain threshold - a function of the number of possible available counterparties - the bank who has to settle the payment is unable to obtain sufficient funds from anyone. The request is cancelled and flagged as defaulted obligation, whose amount is recorded as an instance of insolvency by the short bank, and as a loss by the intended payee bank. Liquidity expectations for both ends of the transaction are adjusted accordingly.

Disaster is simulated through the introduction of an agent of the "disruptive event" breed. This agent can be called from the StarLogo interface and its job is to pick a bank at random and destroy it, turning it into an agent of the "crater" breed. A crater symbolizes a completely inactive bank, who neither makes or receive payment requests, nor operates in the interbank market (as lender or borrower). In the thirty StarLogo seconds after the disruption has taken place, no agent is aware of it, and all surviving banks continue their routine activity. A random process then makes the banks aware of disaster, with the probability of awareness increasing over time. Once a bank is aware, it stops requesting payments from the bank-turned-crater, no longer considers it as a counterparty for loan requests, transforms all payment and loan requests that are pending toward the crater into defaulted obligations, and updates losses and expected liquidity.

Depending on the simulated scenario, the central bank monitoring module may also be activated. Should this be the case, after a given amount of time after the disruption the central bank starts checking whether banks are delaying their payments beyond some pre-determined physiological threshold. If a bank features an excessive number of delays, the central bank supplies it with an amount of cash proportional to them. The intervention routine stops when the number of delays is back below the threshold.

[3] Though not game-theoretically founded, this rule reflects satisfactorily the fact that banks schedule a large share of outgoing payments according to institutional cut-offs. These cut-offs are agreed with customers, who initiate payments, with the receiving bank, when the payment arise from interbank trades (see e.g. the guidelines of the Euro Banking Federation for money market related payments), or established by the system rules (e.g. for payments related to monetary policy operations).

PARAMETERIZATION AND RESULTS

We parameterize our simulation based on real Summer 2007 data for the Italian RGTS system (BI-REL). More than 100 banks participate in BI-REL directly; for the sake of simplicity, we collapse them into five agents. Each agent is a "superbank", incorporating banks that appear homogeneous in terms of payment traffic, opening balances, and collateral.[4]

Table 1 describes the five agents as they are in the real world. Superbanks 1 to 4 are aggregation of Italian banks, while superbank 5 is the aggregation of Italian branches of foreign banks. We draw simulated endowments and payment traffic from random distributions constructed so as to consider:

1) end-of-day liquidity and collateral, used as a proxy of starting values;

2) payment flows per minute among the five entities, obtained by excising certain categories of payments from the original dataset: intra-agent, cross-border, and payments to and from the central bank, whose role in our model is limited to the one of lender of last resort. The quota of neglected payments is subtracted from the endowments of cash and collateral at the beginning of the simulation. The stream of simulated payment flows sticks to the real distribution pattern, highly skewed, with frequent small payments and rare big ones.

Figure 2 depicts the model predictions. The four panels represent the evolution of liquidity levels, liquidity expectations (expressed as differences between current liquidity and liquidity as estimated after the settlement of all payment flows currently existing in the system), money market thickness, and delays incurred by banks in settlement activity.

Under normal operational conditions, the evolution of both liquidity and expected liquidity predictably resembles a random walk; the specific pattern observed in the sample is almost entirely driven by the few large payment orders circulating in the system, and the starting conditions do not appear to generate any path-dependent evolution. Banks rely on the money market infrequently, and mostly when they need considerable amounts of liquidity; the result is consistent with the real share of interbank loans, estimated at 5-7 per cent of the total intraday payment traffic.

TABLE 1 Main scenario *(€ millions)*

Agent number	Average payments settled daily	Large[a] payments settled (% of the total)	Average end-of-day liquidity	Average end-of-day collateral
1	18,384	3.0	5,302	2,648
2	22,669	6.6	3,583	3,947
3	13,718	3.1	4,780	2,766
4	18,451	8.8	2,211	850
5	43,339	7.5	439	10,119
Total	116,562	5.0	16,315	20,329

[a] Payments are defined "large" if their amount exceeds the 95th percentile of the global distribution of flow values.

[4] The reduction in the number of agents, forced by computational limitations, impacts on the evolution of the complex system. The BI-REL system is, however, quite concentrated; two superbanks, for example, correspond to actual single banking groups, with a high level of internal coordination in the payment system.

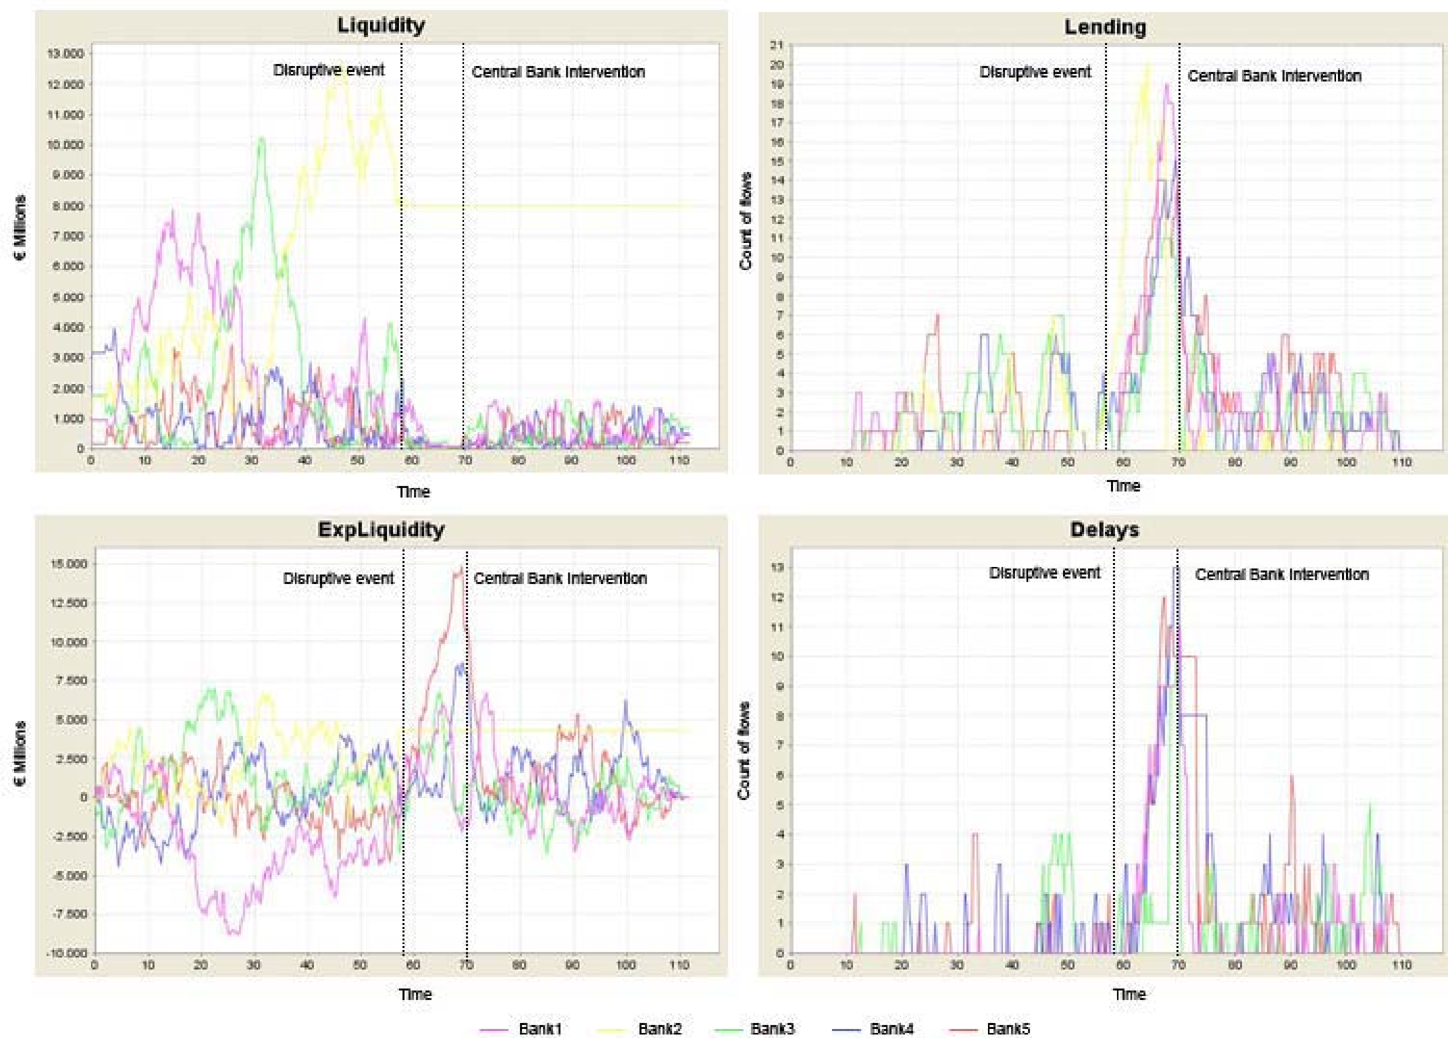

FIGURE 2 Predictions of the model in the minimal intervention scenario

When the disruptive event occurs, and one of the banks is rendered unable to perform any operation, the system evolves through three distinct phases. At first, when no agent is aware of what happened, a spike in expected liquidity emerges. Banks keep on sending payment requests to the bank-turned-crater, incorporating the future settlement of such requests in their expectations; the crater, however, is not able to send out requests of its own, resulting in failure of the regular counterbalancing mechanism for expectations, and illusions of short-run liquidity increases for all its counterparties.

This heralds the second phase, marked by a sizable boom of the money market. Banks experience a lack of liquidity, because they do settle all requests from the crater that were pending at the time of the disaster, but their own requests toward the crater are not settled. After consuming their whole endowment of collateral in exchange for central bank money, they try to counteract the lack of liquidity by turning to the money market. Since loans are granted based on current and expected liquidity both, and expected liquidity is artificially inflated for the reasons stated above, all banks are willing to lend to other banks money they do not yet have: the thickness of the interbank market rises sharply, and as the actual liquidity fails to come in, delays accumulate.

The third phase sets in as banks start to become aware of the disruptive event. One by one, they realize that a bank is not operational anymore, and adjust their liquidity expectations accordingly. Money market activity slows down, and losses are accumulated.[1]

The impact of central bank intervention depends crucially on both timing and intensity. In the scenario depicted in Figure 2 the central bank is relatively slow-moving, and it provides banks with small amounts of extra liquidity: this results in a reduction of delays to physiological levels, but it is not enough to staunch the flow of defaults completely. Runs of the simulation with different behavioral assumptions for the lender of last resort show that the amount of liquidity to funnel so as to neutralize the domino effect entirely can be estimated at somewhere between 1.5 and 2 times the aggregate starting liquidity in the system, depending on circumstances.

CONCLUSIONS AND FURTHER RESEARCH

The model predictions approximate the macro-features of reality adequately, but the framework can be improved along several directions, with the aim to better reflect real RTGS and money market environments. The payment submission process can be refined moving from a direct-debit to a credit-transfer based system, where payments are submitted by the payer. The number of banks should be enlarged and some source of uncertainty could be introduced, by relaxing the independence assumption underlying the payment generation and the common knowledge hypothesis. Moreover, banks should be assigned an end-of-day target in terms of cash balances, to mimic the interday liquidity management optimization they pursue during the maintenance period of required reserves.

As for the central bank, it has to be considered that it autonomously makes and receives payments in real RTGS world, beyond the well-known liquidity supplier function. Modeling of the money market can be refined to take into account the role of overnight interest rates in influencing banks borrowing and lending decisions. Finally, the rest of the world could be

---

[1] Whether the system goes back to normal functioning, aside from the accumulated stock of defaulted obligations, depends on the relative impact of the disruptive event and its consequences compared with the global amount of liquidity in the system. According to our model, the system would not be able to react autonomously, i.e. without bailouts from the central bank, to any crisis neutralizing one of its major players.

introduced at least as an external shocking agent, also to account for the relevant (and increasing) real-life share of cross-border payment traffic.

REFERENCES

A. Alentorn, S. Markose, S. Millard and J. Yang, 2006, “Designing large value payment systems: An agent-based approach”, *mimeo*.

J. Arnold, M. L. Bech, W. E. Beyeler, R. J. Glass, K. Soramäki, 2006, “The Topology of Interbank Payment Flows, Federal Reserve Bank of New York, *staff report* n. 243.

M. Bech, R. Garratt, “The Intraday Liquidity Management Game”, 2003, *Journal of Economic Theory*, **109**: 198-219.

M. Boss, H. Elsinger, M. Summer, S. Thurner, 2004, “Network topology of the interbank market”, *Quantitative Finance*, **4**: 677-684.

G. Fioretti, 2005, “Financial Fragility in a Basic Agent Based Model”, *mimeo*.

M. Galbiati, K. Soramäki, 2007, “A competitive multi-agent model of interbank payment systems”, *mimeo*.

N. Gilbert and P. Terna, 2000, “How to Build and Use Agent-Based Models in Social Science”, *Mind & Society*, **1**: 57-72.

H. Leinonen (ed.) 2005, *Liquidity, risks and speed in payment and settlement systems – a simulation approach*, Bank of Finland Studies, E:31.

S. Wolfram, 1994, *Cellular Automata and Complexity*, Addison Wesley Publishing Company, Reading, MA.