Adapted excerpts from [1], as documentation to ‘An Agent Based Model to assess resilience and efficiency of food supply chains’

Geerten Hengeveld[1], George Van Voorn[1*]

1 Wageningen University & Research, Droevendaalsesteeg 1, P.O. Box 16, 6700 AA, Wageningen, the Netherlands

* Corresponding author. E-mail: george.vanvoorn@wur.nl

Abstract

This is an Agent Based Model of a generic food chain network consisting of stylized individuals representing producers, traders, and consumers. It is developed to: 1/ to describe the dynamically changing disaggregated flows of crop items between these agents, and 2/ to be able to explicitly consider agent behavior. The agents have implicit personal objectives for trading. Resilience and efficiency are quantified by linking these to the fraction of fulfillment of the overall explicit objective to have all consumers meet their food requirement. Different types of network structures in combination with different agent interaction types under different types of stylized shocks can be simulated. Application can be found in [1].

[...]

Agent Based Model

We develop an Agent Based Model (ABM) in order to explicitly include rules of social behaviour and decision-making of agents, the interaction network of agents, and the diversity within the agent population(s) and changes therein, while still being able to include system dynamics [3, 5, 6]. We consider a generic food system, i.e., the model includes main features shared by most if not all food systems, without focusing on the specific characteristics of any particular food system. Note, that in the current model system dynamics (like crop growth and weather processes) have been excluded for simplicity. We realize the generic nature and the exclusion of system dynamics limits the applicability of this model study for specific applications, but again, we are mainly focusing on the trade network structure. Different model settings are explored in a numerical experimental design to generate simulations, which we use to determine what networks are more resilient than others. Moreover, in the context of resilience it has to be clarified “resilience of what to what?” [7]. Hence, several different shocks are included in the numerical design to which the simulated agents can respond. The model coded in NetLogo [8] is available via OpenABM (`https://www.comses.net/`). For the description of the model we loosely use the ODD format for the documentation of ABMs [9].

Purpose of the model

This ABM is intended to represent a generic, dynamic, hierarchical food system network of producers, traders, and consumers. Each model simulation is in fact a trading game that includes a series of production cycles, while in each production cycle several trading cycles are embedded. The ABM is used for studying the trade-off between efficiency and resilience of food supply systems by looking at different network and agent interaction types in relation to different shocks. The food supply chain functioning is evaluated based on the average level of 'satisfaction' of the consumers.

Entities and state variables

The ABM contains a social network in which three types of agents interact. Each agent type represents a different role in the food system:

- Producers, who at the start of each production cycle produce either of two 'crops' (the model is flexible in the number of crops it can handle, but we set it to two by default), and who sell these items to traders based on active selection (see interaction modes further below);
- Traders, who passively accept crops from producers (currency is not explicitly included in the current model code) and passively 'display' crops for selection by consumers;
- Consumers, who actively select traders from whom to buy. Each iteration they try to fulfil their individual dietary needs regarding the existing types of crop.

The main variables per agent type are given in Table 1.

Basic principles, process overview and scheduling

The ABM represents a generalized food system as an abstracted economic game in which crops are produced, traded, and consumed, taking place in a non-spatial, social network of agents. Agents attempt to procure crops in indirect competition. A graphical conceptual overview of the generic model is given in Fig 1.

Several key model assumptions set this ABM apart from most economic models. First, there is no assumption of equilibrium between supply and demand. Food supply is consistently lagging behind food requirement [2], i.e., the minimal nutritional and caloric needs of all consumers will not be met. This is why in this model we set a default value for crop production that is lower than the default value of the requirement. Second, while price formation is an important driver of trade, price is not explicitly included in this model. Instead, in the *weighted* interaction mode (see Subsection on 'Interactions' below), we assume that traders divide their products equally between multiple consumers 'buying' from them, while consumers search for the trader(s) with the largest stock, implicitly assuming that the trader with the highest stock will automatically offer the lowest price. Vice versa, producers look for traders with lowest stocks with the assumption they will offer higher prices. In other words, agents behave according to economic utility, but in the ABM the relative availability of crops is taken as proxy for the relative price traders ask. No explicit monetary stocks and flows are modelled. Untraded crops remain in stock for possible trade in consecutive iterations, although the parameter setting can be changed to include a loss term from storage. These assumptions present an important reason why the ABM is a dynamic model instead of an equilibrium model.

The order of actions by the different agents is as follows. Each production cycle ('year') all producers choose a crop, and produce it. Each producer chooses a trader to sell their produced crop to (based on indirect economic utility, and possibly modified, see the below Subsection on 'Interactions'). Each consumption cycle (interpreted as 'month' when the parameter for the number of consumption cycles is set to '12') trading takes place. Each consumer chooses a trader to buy from (again based on indirect

Table 1. Main variables per agent type.

Agent type	Variable	Description	Type	Temporal dynamics
Producer	Crop_type	The crop type currently produced	Int	Each iteration
	Production	The volume of the crop currently produced	Double	Changes during specific shock period
	Production memory	The volume of crop produced at the start (to be able to get back after shock)	Double	Fixed
	My_trader	Trader to whom current production is sold	Agent	Each trade iteration
	My_historictraders	List of traders to whom crops have been sold in the past	Agent set	Accumulates each trade iteration (no loss of memory)
Trader	Producerset	Current producers from whom is bought	Agent set	Each trade iteration
	Consumerset	Current consumers to whom is sold	Agent set	Each trade iteration
	Product quantitylist	Amount of crop in store per crop type	Double [NCroptypes]	Each trade iteration
Consumer	Consumption QuantityList	Amount of crop consumed per croptype per period this tick	Double [nPeriods,NCroptypes]	nPeriods per Tick
	Demandlist	Amount of crop required per period per croptype	Double [nPeriods,NCroptypes]	Changes with specific shock period
	Demandlistmemory	To be able to recover after a shock	Double [nPeriods,NCroptypes]	Fixed
	MyTrader	Trader currently bought from	Agent	nPeriods per Tick
	MyHistoricTraders	List of traders from whom crops have been bought in the past	Agent set	Accumulates

Computer program variables used by the different agent types.

economic utility). Trading interactions are direct. We assume a strictly hierarchical, three-level layered network structure, in which the agents are nodes, and the trading with crops forms the unidirectional flow links (again, the money flow in the reverse direction is in this model formulation only implicit), in which each link represents a connection from agent i to j, and has time-dependent attributes like capacity and flow rate [4]. For our analysis we assume that agents share a link if they (commonly) trade during some time interval – and, hence, no link exists if within that time interval no crops are exchanged from agent i to j. Producers only trade with traders, who in turn only trade with consumers, and no direct trade between producer and consumer can take place. Traders are therefore a necessary intermediate in the flow of crops from producers to consumers. Traders split the crops they have in stock evenly between the different consumers buying from them, up to the requirement of these consumers. Left-over crops are stocked for the next consumption cycle. The trading network is organized according to one of several pre-defined types (discussed further below). Trading takes place according to one of three possible interaction modes: random, according to basic economic utility optimization, or a limited economic utility optimization, constrained by preference for known traders (also discussed further below).

Further important model assumptions are:

- Links between individual agents can change in time according to one of the three (below-described) interaction types, i.e., links between agents can (dis)appear during a simulation;
- The number of agents of each type remains fixed during a simulation, and agents do not switch agent type, i.e., they keep their place in the agent network. Agents cannot disappear from a simulation, even if they would go 'bankrupt' or do not fulfil their dietary needs (this assumption is justified here as we are interested in the impact of shocks, and not in the exact development of the food system in time), and no new agents are introduced during a simulation.

Simulations take 300 iterations, of which the first 100 present a 'start-up' phase, and the next 100 the 'pre-shock' phase. After 200 iterations a shock period takes place in which one of three possible shocks occurs that lasts 100 iterations. The severity of the shock is determined as a severity parameter (set to 0.5 by default) times the number of randomly selected agents of the level at which the shock is applied to. The three potential types of shock are:

- Shock at the producer level, which could e.g. be interpreted as 'failed harvest'. A fixed percentage of producers – determined by the severity parameter – will not produce crops. This reduces the total volume of crops available to traders to buy;
- Shock at the trade level, which could e.g. be interpreted as 'stock loss'. A fixed percentage of traders – determined by the severity parameter – will lose their complete stock. This reduces the total volume of crops available to traders to sell;
- Shock at the consumption level, which could, e.g., be interpreted as 'preference change'. A fixed percentage of consumers – determined by the severity parameter – will shift their preference from one crop type to another. This changes their objective to fulfil their individual requirement, and changes the requirement of consumers of what to buy from traders.

Scales

The ABM does not have an explicit spatial scale. Although agents are shown at specific locations in the simulation environment, this does not mean there are explicit spatial processes. Instead, there is only a social network of producers, traders, and consumers, where agents have a 'position' in the network without an explicit spatial location. The implicit assumption following from this is there are no costs for transport. Although the agent typology implies agents to be individuals, agents can also be interpreted as companies, super-individuals, or even countries in their roles of producing / exporting and consuming / importing crops. The temporal scale is set to 1 production cycle and 12 consumption cycles ('months') per tick, mimicking a yearly cycle with a single production of crops each year, in which consumers can change supplier each month. The difference in temporal scales represents the higher flexibility of consumers for choosing where and when to buy, then for producers when to sell. In the analysis the effects of stochasticity are dampened by the use of the average over a large numbers of iterations.

Interactions

There are three interaction modes one can select from when initializing the model (i.e., once the interaction mode has been set, it keeps this mode during simulation):

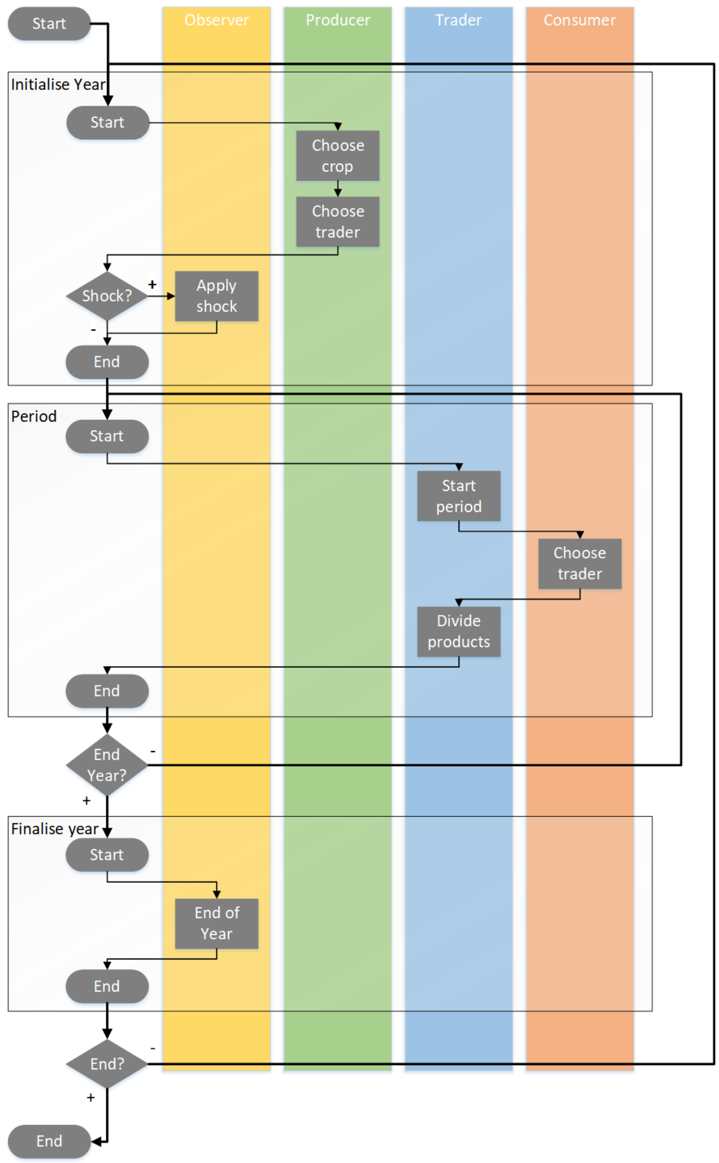

Fig 1. Flowchart of the model. The ABM consist of a level of producers (in green), traders (in blue), and consumers (in coral); 'observer' (light orange) deals with user-defined decisions and actions. Agents select partners to trade with. Traders also have the ability to stock to sell later. 'Year' refers to production cycle. 'Period' refers to consumption cycle.

- The *random* mode. This entails the random selection of a partner. There are no explicit decisions but only random pairing of producer with trader and trader with consumer. This also means trading is **not** according to economic principles. This mode presents a 'baseline';

- The *weighted* mode. Selection occurs according to 'best interest', i.e., agents evaluate which of the possible trading partners offers the best price for buying or selling, respectively. As mentioned above, price optimization is implicit. Producers select traders negatively weighted by the traders' stock of the crop the producer is producing, mimicking the highest selling price. Consumers select traders positively weighted by the traders' stock of the crop for which the cumulative fulfilment of the requirement is lowest, mimicking the lowest buying price. Selection is based on perfect information, which is a common – yet admittedly doubtful – assumption in economic theory. Excess crop volume is stocked by traders;

- The *preference* mode. Selection is based on a combination of perceived best price and a preference based on trading history. This is hence similar as for the weighted mode, but now those traders with whom has been traded in previous production and consumption cycles get additional weight in the selection process. This interaction mode mimics price sensitive behaviour constrained by preference for trusted partners. This is a reasonable assumption, as trust plays an important role in trading.

Also, one can select from seven possible networks (see Table 2), divided over four different types:

- Block type, i.e., there are equal numbers of producers, traders, and consumers;

- Inverse pyramid type, i.e., there are few producers, more traders than producers, and more consumers than traders;

- Hourglass type, i.e., there are more producers and consumers than traders;

- Diamond type, i.e., there are more traders than producers and consumers.

All three interaction types and all seven networks have been included in the numerical experimental design (see next Subsection).

Table 2. Network types included in the numerical experiment.

Network type	Network code	No. of producers	No. of traders	No. of consumers	No. of possible links
Block (small)	bs	5	5	5	50
Block (large)	bl	20	20	20	800
Inverse pyramid (small)	is	5	25	50	1375
Inverse pyramid (large)	il	10	50	100	5500
Hourglass	h	25	5	50	375
Diamond (small)	ds	5	50	25	1500
Diamond (large)	dl	10	100	50	6000

Columns including abbreviations (network code), number of producers, traders, and consumers, and the number of possible links.

Objectives

We distinguish between two types of objectives. All agent types (and hence all individual agents) have *implicit personal* objectives when trading. They (implicitly) aim to buy or sell at the best price (or a combination of best price and preference, explained further in the next Subsection). This applies to producers, traders, and consumers.

Note, that the price dynamics are not explicitly included in this model, but we make the assumption that 'price' is determined implicitly by the size of the stock the trader has available. Producers sell off their product to traders who offer the best price, which is based on the trader having low stocks, which implies he offers a good price. Consumers buy at the lowest price. This is based on the trader having high stocks, which implies he is willing to accept a lower price. Traders do not 'need' to sell, and may 'choose' to store crop items for later transactions. These (implicit) 'decisions' are the result of a failure to link to a consumer in a certain iteration. The *explicit system-level* objective we consider in our analysis is the fulfilment of nutritional requirements, and in line with how we define efficiency (the share of produced food delivered to consumers). Consumers have the objective to fulfil their individual nutritional and food preference requirement, which is translated as a requirement that is set as a parameter for each crop type each cycle (while, as said before, they implicitly attempt to do so at the best price, or a combination of best price and preference). If they do not manage to fulfil these demands, their satisfaction drops (note, that in the model we now assume there is no starvation, bankruptcy, or any other adverse affects from failure to fulfil the requirement). During simulations the food system performance is evaluated through the fulfilment of consumer dietary needs: consumer satisfaction is measured as a discrepancy between what a consumer needs of each crop type each production cycle and what he manages to obtain via trading. Note, that because there is an inherent imbalance between supply and demand (i.e., shortage) in the model, on average consumers will not achieve 100 % satisfaction (but individual consumers may do so). The satisfaction is calculated as the minimum over the crop types of the ratio of consumption to requirement (explained further in the next Subsection).

Emergence

The model is used to study the quantification of resilience to shocks in food system networks. The networks are formed through the exchange of crops. Although the number of agents in a simulation is 'fixed' and producers cannot directly trade with consumers, the network of exactly who trades with whom emerges through these exchanges of crops, depending on the interaction rules. The resilience of the food system emerges from the capability of agents to change with whom they trade, and the capacity to take stock. Resilience is assessed as the difference in flows of crop volume between the shock and pre-shock period (explained further in the next Subsection). These flows are approximated by averaging the total crop volume traded between two agents over one of these two periods. Average flows differ between model configurations, determined by network type, interaction type, and shock type. Network types determine the 'hard' boundaries for the capacity to redirect flows. Interaction types determine the 'soft' boundaries, for instance, some potential links may exist but are never created due to preference interactions. The emergent property that is modelled is the (relative change in) resilience of the different food system network types and interaction modes in response to different shocks.

Adaptation

The main source of adaptivity in this model is the capacity of agents to trade with other agents. Rules for trading depend on the interaction type. With random interaction selection the adaptation comes from the random making and breaking of trading links. The resilience is determined by the 'hard' boundaries, as in principle all potential links can be made (note also, that because the pairing is random trading is usually not optimized, and there is no penalty for poor fitness in this model, as agents cannot disappear from the simulation). If a shock results in the cutting of links, the random pairing 'fixes' this. With non-random interactions there are 'soft' boundaries to the generation of trading links. Links that would be feasible in theory (with random pairing) are now practically infeasible, which implies a reduced adaptation capacity. This is particularly the case for the preference interaction mode.

Learning

There is no specific learning, other than that agents 'remember' with whom they have traded and assign weights to that in case of preference interactions.

Prediction

Agents predict optimal trades in case of weighted or preference interactions.

Sensing

Agents have perfect information on crop availability and 'sense' what traders have available. Agents also 'sense' their own internal states (the available crop, their weights to other agents). Interactions include 'full' information disclosure about these states to (potential) trading partners.

Stochasticity

The model incorporates two sources of randomness:

- Each production cycle the producers are assigned a random crop type to produce;
- The selection of traders by producers and consumers is a random or a weighted random process.

To avoid dominance of model behaviour by stochasticity all other parameters, variables and processes are approximated to be fixed for the entire population or period.

Initialization

The model includes 16 parameters that need to be set at initialization. The use of some of these is conditional on the value of others. These parameters are shown in Table 3. At initialization all agents are placed in the field. Producers are assigned a fixed production. Consumers are assigned a crop type each consumption cycle, depending on what is available.

Observation

Output is discussed in detail in the next Subsection. The simulation runs are saved and all data from it are available, i.e., no data limitations or added noise have been considered to mimic real-life data properties.

Input data

The model represents an iconic, generalized food system that is not calibrated to a particular real-life food system. No specific input data is used.

Table 3. Parameters of the model.

Parameter	Description	Default value	Units
N_Producers	The number of producers	$[1, \rightarrow\rangle$	(Agents)
N_Traders	The number of traders	$[1, \rightarrow\rangle$	(Agents)
N_Consumers	The number of consumers	$[1, \rightarrow\rangle$	(Agents)
TotalProduction	The total volume of crops produced by all producers combined per production cycle	100	Volume
TotalDemand	The total requirement by consumers for all crops combined per production cycle	120	Volume
N_CropTypes	The number of crop types in the simulation	2	–
StockPersistence	Percentage of stocked crops that persist into the next production cycle	0.95	Tick^{-1}
Periods	Number of consumption cycles per production cycle	12	Tick^{-1}
SpinupPeriod	Number of production cycles taken for spin up of the model (no output produced)	100	Ticks
StationaryPeriod	Number of production cycles taken as the standard situation	100	Ticks
ShockDuration	Number of production cycles the shock is applied	100	Ticks
ChoiceModel	The used interaction mode	Random / Weighted / Preference	–
BasePreference	Value used as basic weight for all traders for weighted and preference interaction mode	0.01	–
ExpPreference	Value used as weight for known traders in preference interaction mode	100	–
ShockType	Type of shock	Producer / Trader / Consumer	
ShockSeverity	Severity of the shock / share of population of targeted agents affected	0.5	–

The parameters of the model with description, default values and units. Parameters included in the sensitivity analysis are indicated in grey. Time units are set in Ticks, the default time setting in NetLogo, with the arbitrary interpretation of a year. The parameters included in the Sensitivity Analysis are indicated in grey.

Output data

The model output is organised such that high-resolution insight can be gained from the trade interactions between the different agents for a large number of simulations, e.g. run through behaviorspace. The output consists of:

metadata.csv lists the settings for the different runs executed, including the number of agents, configuration and the run identifyer linking each run the the run specific output files.

outputfile_*i*_*j*_general.txt the per tick summary of the state of the food system. Reports state of stocks and fulfillment of food requirement (health) aggregated over the agents, and for the resources both aggregated and separately. This file is specific to run j of behaviorspace experiment i.

outputfile_*i*_*j*_linkmatrixheader.txt provides the run specific header for the link matrix output files. This file is specific to run j of behaviorspace experiment i.

outputfile_*i*_*j*_linkmatrix_all.txt provides the per tick tradelink matrix aggregated for all resources. This file is specific to run j of behaviorspace experiment i and needs the associated linkmatrixheader file to interpret.

outputfile_*i*_*j*_linkmatrix_*k*.txt provides the per tick tradelink matrix for resource k. This file is specific to run j of behaviorspace experiment i and needs the associated linkmatrixheader file to interpret.

The ../results/ folder should exist to generate the output files.

Model analysis and quantifying resilience

The calculation of the resilience and efficiency metrics by [1] are post-processed calculations. In [1], the model is analyzed by using a numerical study design involving the three qualitative input factors (interaction mode, network, and shock type) and an OFAT (One Factor at A Time) Sensitivity Analysis involving continuous parameters. Each involved simulation consists of 300 iterations: 100 iterations for initialization, followed by 100 iterations pre-shock period, and then 100 iterations with the system under shock. The simulations presented in [1] can be recalculated using the different behavior space experiments involving 8100 runs. Take note, output files are of considerable size. Only a subset of the main experiment is included in the analysis in [1].

[...]

Acknowledgments

This research was funded by the Strategic Investment Programme on Resilience of Wageningen University & Research and the KB 35 “Food Security and Valuing Water programme” that is supported by the Dutch Ministry of Agriculture, Nature and Food Security. The funders had no role in study design, data collection and analysis, decision (what) to publish, or preparation of the manuscript.

References

1. Voorn GAK van, Hengeveld GM, Verhagen J. An Agent Based Model representation to assess resilience and efficiency of food supply chains. Plos one. *in press*
2. Ray DK, Mueller ND, West PC, Foley JA. Yield Trends Are Insufficient to Double Global Crop Production by 2050. PloS one. 2013;8(6):e66428. `https://doi.org/10.1371/journal.pone.0066428`
3. Gotts NM, van Voorn GA, Polhill JG, de Jong E, Edmonds B, Hofstede GJ, Meyer R. Agent-based modelling of socio-ecological systems: Models, projects and ontologies. Ecological Complexity. 2019;40:100728. `https://doi.org/10.1016/j.ecocom.2018.07.007`
4. Levalle RR., Nof SY. Resilience in supply networks: Definition, dimensions, and levels. Annual Reviews in Control, 2007:43;224–236. `https://doi.org/10.1016/j.arcontrol.2017.02.003`
5. Macal CM. Everything you need to know about agent-based modelling and simulation. Journal of Simulation. 2016;10(2):144–56. `https://doi.org/10.1057/jos.2016.7`
6. Schulze J, Müller B, Groeneveld J, Grimm V. Agent-based modelling of social-ecological systems: achievements, challenges, and a way forward. Journal of Artificial Societies and Social Simulation. 2017;20(2). `http://jasss.soc.surrey.ac.uk/20/2/8.html`

7. Carpenter S, Walker B, Anderies JM, Abel N. From metaphor to measurement: resilience of what to what?. Ecosystems. 2001;4(8):765–781. https://doi.org/10.1007/s10021-001-0045-9

8. Wilensky, U. NetLogo. https://ccl.northwestern.edu/netlogo/

9. Grimm V, Berger U, DeAngelis DL, Polhill JG, Giske J, Railsback SF. The ODD protocol: a review and first update Ecological modelling. 2010;221(23):2760–8. https://doi.org/10.1016/j.ecolmodel.2010.08.019

10. ten Broeke GA, van Voorn GA, Ligtenberg A. Which sensitivity analysis method should I use for my agent-based model?. Journal of Artificial Societies and Social Simulation. 2016;19.1. https://doi.org/10.18564/jasss.2857

11. Ulanowicz RE, Goerner SJ, Lietaer B, Gomez R. Quantifying sustainability: resilience, efficiency and the return of information theory. Ecological complexity. 2009;6(1):27–36. https://doi.org/10.1016/j.ecocom.2008.10.005