

Open- and Closed-Loop Supply Chain Dynamics: Specification and Exploration of an Agent-Based Model¹

Christopher S. Ruebeck, Department of Economics, Lafayette College
Jeffrey O. Pfaffmann, Department of Computer Science, Lafayette College

Abstract

We develop an inter-firm shipping pallet supply chain model and explore the geographic features of market structures that can evolve within it, focusing on the common resource nature of pallet durability. While exploring the model's parameter space we find regions in which the outcomes do not put the model's assumptions in doubt. The insights reveal modeling considerations for an "open-loop" structure, where the shipping pallets received by retailers pass through third parties before returning to manufacturers. Our approach to validation is at the level of "stylized facts," accepting model parameters that produce outcomes matching characteristics of firms' geographic location.

Introduction

Environmental concerns are important drivers of innovation today, and such issues are pervasively linked to the supply chain of every good we produce and consume (Thomas and Griffin 1996). The supply chain itself has impacts that are difficult to model (Mabert and Venkataramanan 1998). We study an important component of delivering goods from producer to retailer: the shipping pallet.

Shipping pallets are used at every stage of production, including distributing material within the firm. For simplicity, we focus on the final step of shipment to the retailer, a significant component of total pallet use (Freedonia, 2008). This final step is particularly important because the pallet is moving between economic actors. While a pallet stays within a firm, the firm's incentives to manage its private resources align well with the optimal social outcome, but as the pallet moves between firms, natural externalities may cause the firm's optimal choice to diverge from social welfare.

This investigation is part of a larger effort to include the effects of private incentives—market interactions—in environmental policy considerations, particularly when performing a life cycle analysis (LCA) on a system of interest. We are not the first to use agent-based modeling in a supply chain context (some examples: Swaminathan et al. 1998, Choi et al. 2001, Gosling 2003, Hensher and Puckett 2005, Gosling et al 2006, Chaturvedi et al. 2006). Our innovation is to link

¹ We thank Sharon Jones for her contributions to this model's development; Nicolaas Vriend and attendees to the 15th International Conference on Computing in Economics and Finance for helpful comments; and Rachel Ingber, Chinh Do, Ed Seo, Anthony Finizio, Aung Nang Lin, Jonathon Jenkins and Brian Warner for research assistance. We are grateful for funding through NSF grants HSD BCS-079458 and CPATH-T 0722211/0722203.

this analysis with policy making. To date the LCA policy tool does too little to recognize the interaction between policy and markets (Fischhoff and Small 2000) and also does not address durable goods issues well (Cooper 2005). Our approach to melding policy interests and markets is through agent-based modeling of the market itself. This will allow the policy maker to observe and interact with the model as it demonstrates the consequences of changed market incentives.

In this paper we describe our model's structure and the results from validating at the level of a "stylized fact". In particular, we describe valid parameter ranges as those that reproduce the observed market behavior of repair firms' tendency to cluster near pallet users. This validation method is necessary when data on participant behavior is difficult to attain, and recognizes that any market model must be an abstraction of the true system's complex details. We can thus incorporate both market parameters that we have observed and abstractions of the market structure that keep the model tractable.

Our model also abstracts significantly from reality, maintaining an appropriate level of complexity for an initial investigation into a model of this type along with its method of validation. Validation methods for agent-based modeling are in their infancy, and models currently developed often use agents of very few types (Lux 2009, Franke 2009). We discuss the additional complexity that can be modeled with later efforts in a section below on "the big picture" and in our concluding section.

We next describe the market and its relationship to our model's structure, followed by our methodology and the experiments we have used to investigate the model's functionality, and we finish with a description of future work.

Background: the market for shipping pallets

Basics

Our model focuses on the market for used shipping pallets so that we may explore the relationship between open-loop, closed-loop, mixed, and hybrid systems for managing the pallets in an inter-firm supply chain. The production of new pallets acts as a source to this system, and pallet disposal by landfill or mulching acts as a sink. We further constrain the focus of the model to wood pallets, for multiple reasons. The most obvious is that wood pallets were 80% of new pallets in 2007 and are predicted to continue to be the "vast majority" of pallets in the future (Freedonia, 2008). For our effort, though, the most important reason to focus on wood pallets is that this pallet type is both more amenable to repair and less durable than other pallet types. Additionally, and related to both of those reasons, other types of pallets are primarily found in closed-loop relationships, while wood pallets are found in all supply chain relationships. Thus if we want to understand inter-firm effects on the environment through shipping pallets, wood pallets are where the action is.

As part of shipping goods downstream towards the consumer, a purely closed-loop system returns the pallets directly to the source of the goods that were shipped on the pallets; the pallets do not leave the boundaries of a single producer. Purely closed-loop systems are rarely found

outside vertically integrated firms, but there are often relationships across firm boundaries that are nearly closed-loop because of their long-term nature. The U.S. Post Office and its graphics and printing customers have such a relationship, in which pallet ownership remains with the USPS. In a purely open-loop system, the upstream (producing) and downstream (retailing) entities are not only no longer part of the same firm, but relationships between the upstream and downstream entities are also more fluid, and third parties manage the pallets' return upstream. A large fraction of pallets in the inter-firm supply chain are part of these open-loop "systems", reflecting the decentralized nature of a capitalist economy.

Both mixed and hybrid forms of open- and closed-loop systems also exist. In a mixed form, the producer that has integrated forward into retail may receive products from outside sources at its retail locations, and it may also have franchisee retail locations. Thus both producer and retailer pallet users may send and receive pallets through both closed-loop and open-loop systems. Even if not under common ownership, a producer and a retailer may contract to have pallets returned directly to the producers, but these contracts may not exist with all the entities with which either side does business.

Hybrid forms of pallet management come in several varieties. Cooperatives can be formed by geographically-localized retailers that receive shipments from a common set of producers, decreasing costs for both groups. As warehouse outsourcing has grown over the past decade, so have pallet-management companies, the largest of which is CHEP. These firms have actual ownership of the pallets, reinforced by legal findings (Brindley, 2009). Although a significant part of the market and growing, these companies are still small in comparison to the entire market. Our model does not yet represented these managed solutions, instead focusing on the decentralized form, and flexibly designed so that we may add the additional layer in the future.

Contribution to a "common resource"

Standard-size multi-use pallets are thus most often an example of a *common-resource* good when they are in an open-loop system because they are rival but not excludable. Pallets used by manufacturers are rival because using them to ship a good decreases the benefit that anyone else can receive when using the pallet for future shipping. Yet once the pallet is shipped to a retailer, with no contract to enforce its return to the manufacturer, the further use of that pallet is non-excludable. On the other hand, specialized pallets, as well as any pallet used in a closed-loop system or owned by a pallet management company, can be cast as *private* goods because the implicit or explicit contracts between the entities described above make the future use of the pallets excludable as well as rival (Coase 1960). We discuss these pallet management companies further below.

We can thus cast the manufacturer's pallet purchase decision as a contribution to a common resource when in an open-loop relationship. The new pallet itself is not a common resource; rather, future uses of the pallet are the common resource. So more precisely, the choice of pallet durability is the common resource problem. Thus our investigation is one of economics and policy rather than simply supply chain analysis.

The bigger picture

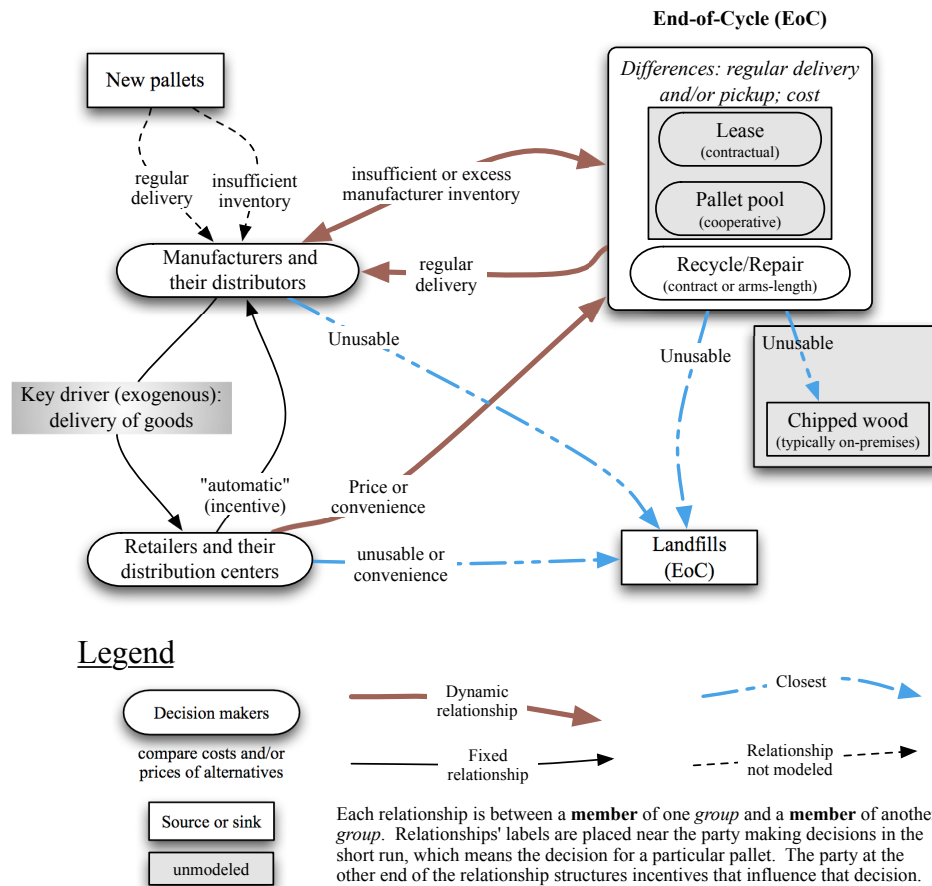


Figure 1. The relationships in our model, including additional end-of-cycle behavior not yet modeled.

To give context to our focus on open- and closed-loop structures in used pallet market, we now describe a more complex “hybrid” situation, one not yet part of the model. This example reflects the large range of firm sizes for each firm type—pallet manufacturer, pallet repairer, and pallet pooler—and the large number of firms that perform more than one of these functions. Figure 1 captures both the simple more complex models. The focus outside this section is on the smaller, local repairers that can enter and exit the market easily and frequently. This richer set of examples we present now highlights the incentives that all market participants face.

Some firms produce a specialized pallet for a particular industry or a particular manufacturer and take on every stage in the pallet’s life: assembly, delivery to the upstream firm, pickup from the downstream firm, repair, and return to the upstream firm. Transportation may be subcontracted. A label on the pallets with a phone number specifies that the pallets will be picked up by the pallet firm (or its subcontracted hauler) within a maximum number of days after calling in for removal, and within an additional number of days the pallet firm is contractually bound to return the pallets, repaired if necessary, back to the upstream manufacturer for future goods deliveries downstream. This turnaround time is on the order of a week.

We can learn about the incentives in the supply chain of *standard-size* pallets by thinking about each step of the *specialized* pallet’s supply chain. The incentives are obvious at each end of the cycle: the upstream goods firm needs pallets and the pallet firm needs to sell them. It is in the middle of the loop that incentives are most important and where aligning those incentives may be difficult.

First consider the downstream goods firm, perhaps a retailer (although either the upstream or downstream firm may be an intermediate goods producer). The retailer's warehouse area can be efficiently cleared by calling the phone number on the pallets; this becomes part of the retailer's routine. On the one hand, if that phone number were not available then the retailer would have to arrange for clearing its loading area of the specialized pallets that have little use to other upstream goods firms. In addition, the contract that the retailer enters to receive the goods includes a penalty if the phone number on the pallets isn't called.

Return now to the upstream goods manufacturer: What is its incentive to create such a contract, one that can include many downstream firms and does not need enforcing frequently? The contract is in this manufacturer's interest if it leads to a lower effective pallet cost.

Finally we come to the pallet firm that has established itself in a market niche with a differentiated product, an effective strategy if it is too expensive for a second pallet firm to enter. This differentiated product is the key driver in such a hybrid (mix of open- and closed-loop) structure.

Specialized pallets are about half the market, and standard-sized "grocery store" 48"x40" pallets make up most of the other half. Thus there are many firms working with standard-size pallets. The firms following the hybrid form may instead outsource their "pallet services" to pallet management firms, for either standard-size or specialized pallets. The data (Freedonia, 2008) tell us that these outsourced operations have become increasingly viable over the past 20 years and are expected to continue in the future. The most well-known firm that provides pallet services is CHEP, taking the hybrid approach further still: this and other 3rd-party firms own the pallets. Although they are big, they each only comprise a little more than 10% of the pallet management market. We will discuss these firms again when we consider agent interactions in our model. Although outsourced pallet management is an important component of the industry, much of the market and its environmental impact remains in the more de-centralized repair firms' hands.

For the remainder of the paper, we only consider standard-sized pallets and small firms that can easily enter and exit the market.

Background: abstraction and methods

Our investigation focuses on three aspects: market reactions embodied in the model, policy alternatives specifying different parameter settings, and computational exploration of parameter settings not specified by policy alternatives. The computational explorations provide an automated mechanism for examining the interplay between embodied market reactions and parameterized policy alternatives.

One challenge with using agent-based modeling to simulate markets is finding the correct validation level: at one level agents can have a high fidelity to the actual decision making process, while at a different level, decisions are represented more abstractly (Miller and Page 2007, North and Macal 2007). An intuitive argument is to interview participants to develop a well established decision making process such as in wholesale electric power markets (Koesrindartoto et al. 2005, Macal and North 2005) or labor markets (Chaturvedi et al. 2005),

but the literature indicates this approach can be problematic. Although the rules by which decision makers operate may be clearly defined, then we know what questions to ask the decision makers and they know what kinds of answers we expect. This approach has two faults: first, limiting the system to a set of answers suggests that the market is perfectly-understood, and second, decision maker interactions in most markets are not nearly so well-delineated (as they may be, for example, in electric power).

Our model focuses on the abstract representation of the inter-firm supply chain with the goal of enabling a user to gain insights while exploring possible scenarios. Thus, the model has many parameters to manipulate, but seeks to provide a balance between parameter-complexity and the ability to describe the market flexibly. To help the user explore the model, a set of automated tools can illuminate where a concept may have been missed. Inspiring this approach is Miller's (1998) technique: a model without surprises is a model that can't be "broken" is also a model that has little meaning.

... these [automated] searches are a means by which to uncover potential weaknesses in the model's formulation and identify key assumptions. With this information, the model can either be refined or, if it is felt to be sound, additional effort can be focused on better estimating and understanding the behavior of the key assumptions. Note that the inability to "break" a model in this way does not guarantee its quality. For example, models that are completely insensitive to their parameters can obviously not be broken in the above manner—yet, such models are also not likely to be of much value. The tradeoff between the brittleness in a models and its responsiveness to parameters will always need to be carefully considered.

We return to a discussion of this approach after describing the model and our experiments.

The explorations in this paper use relatively homogeneous agents. We can choose to vary the choices that agents make, but presently the structure of the agents' decisions are uniformly applied. We will introduce more heterogeneity into the agent populations once we have explored the system behavior under a homogeneous population. This experience will help us to understand the manner in which to specify the heterogeneity using the agent behaviors that we have already encoded, to incorporate hybrid inter-firm pallet supply chains, and to have a solid basis for exploring policy alternatives.

What follows next is a detailed description of the model: agent description, agent interactions, and experiments confirming some ranges of desired model behaviors. Appendix A lists the main loop of code.

The model

To abstract from the market to our model, we conceive of two types of agents: those that use pallets and those that process the pallets at the end of a shipping cycle. In a closed-loop system the two are part of the same firm, but in an open-loop system they are separate. Figure 1 still applies to this description.

The pallet users in our model are goods manufacturers and goods retailers. Manufacturers make more decisions than retailers: they bring new pallets into the system, they can hold pallet inventory across periods, and they buy used pallets from repairers. Both of these users send pallets to “end-of-cycle” firms, the repairer firms and landfills.

Repair firms are the focus of our model; they are the firms that return the pallets to manufacturers in open-loop situations. Landfills are end-of-cycle firms, but would be better labeled end-of-life: they receive and dispose of pallets from manufacturers, retailers, and repairers.

The model has short run, medium run, and long run timeframes. In the short run, pallet users make decisions to buy or sell, and whom to buy from or sell to. In the medium run, repair firms consider exit, and new repair firms enter. Repair firms’ entry and exit are the evolutionary pressures in our model. In the long run we can observe the model’s properties. For example: How do the differences between goods manufacturers’ decisions and retailers’ decisions affect the differential survival of repair firms near the two types of pallet users?

The simulation itself can be run with or without a graphical interface. The mode with graphical interface active is for a single user to explore parameter configurations. The mode without the graphical interface is used for automated search of the parameter space by a specified algorithm. The following descriptions of the global variables, agent variables, and interactions between agent types focus on the graphical interface, but much of this information also translates into the non-graphical operating mode.

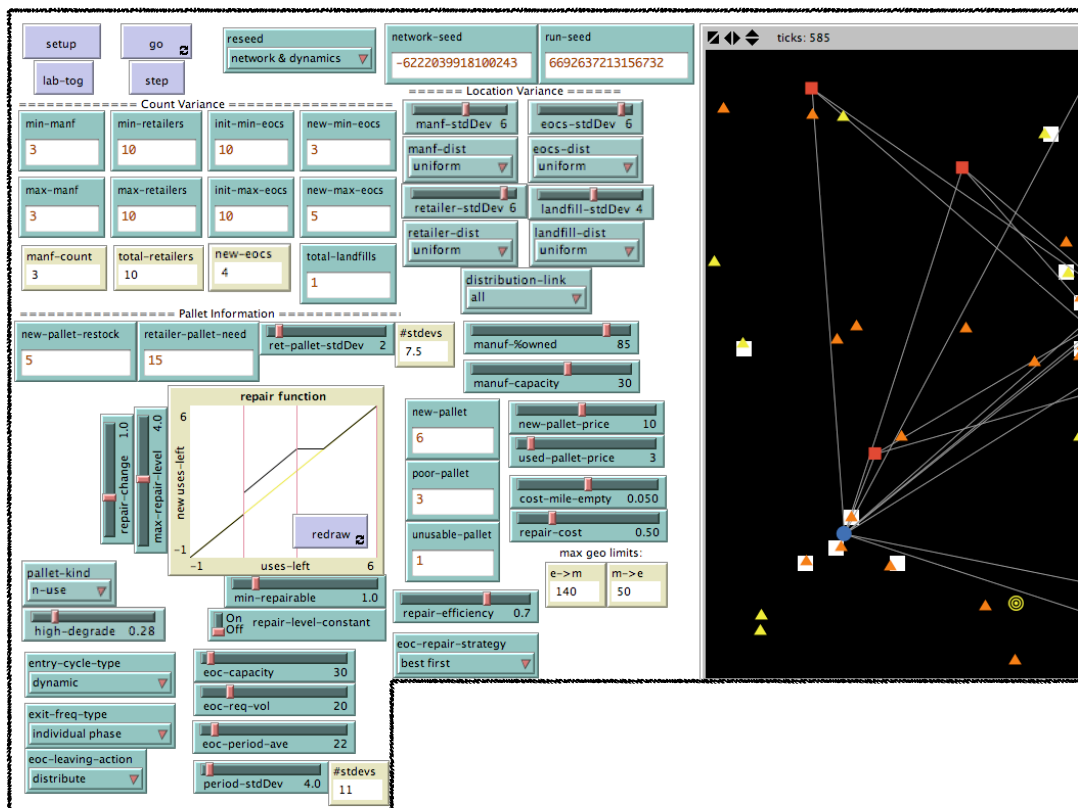


Figure 2. The model in interactive mode, showing the various controls.

Parameter specification

The discussion below groups the model parameters and controls into these categories:

- simulation controls,
- agent population organization,
- agent behavior dynamics, and
- pallet dynamics.

Figure 2 first shows the the interactive mode as a whole, The dark (green) elements define global variables' values and the light (tan) elements display values that are calculated dynamically in the code.

The simulation controls shown in Figure 3 allow the model to be initialized (*setup*) with or without labels (*lab-tog*), run until paused (*go*), and run one time step (*step*). The model uses two different random seed values to separate the initial spatial configuration from the model's execution. This allows users to start with a common spatial configuration but different event sequences, or start with different spatial configurations for every run. Each run in the experiment results reported below uses a randomly chose spatial configuration, with two different random numbers for each run.




Figure 3. Simulation controls.

Figure 4 contains the agent population variables that specify the supply chain infrastructure, the environment within which our agents operate. This environment consists of static pallet users (manufacturers and retailers) and landfills, whose locations and relationships are established when the model is initialized. The links from retailers can either be to all manufacturers or the closest one. The dynamic third-party repair firms (End-of-Cycle agents, or EoCs) do not have location specified here; they enter and leave the environment based on individual decision making behavior.

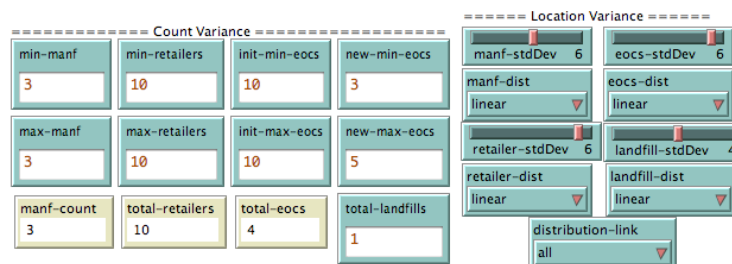


Figure 4. Agent population parameters.

When the simulation is initialized, the number of manufactures, *manf-count*, is chosen from a uniform distribution over the range [*min-manf*, *max-manf*]. Manufacturers can be either uniformly distributed in the space or normally distributed with mean at the world's center and *manf-stdDev* standard deviation. Initial numbers and locations of retailers (*total-retailers*), landfills (*total-landfills*), and repair firms (*total-eocs*) are determined in a similar fashion,

with the added option that they may be clustered in a normal distribution around the manufacturers. The number of repair firms is generated initially from the range [init-min-eocs, init-max-eocs], and during model execution from [new-min-eocs, new-max-eocs]. The timing of EoC entry is described below. Finally, distribution-link specifies that retailers either be linked to (do business with) all manufacturers or just to the closest one.

One set of agent behavior dynamics, whose parameters are displayed in Figure 5, determines the scale of pallet needs between users and the users' contractual relationships. Manufacturers receive a restocking of new pallets each period; goods are shipped on these, on used pallets from inventory (up to manuf-capacity), and on any additional new pallets purchased to satisfy a stochastic need from retailers. The retailers' need is determined separately for each retailer-manufacturer link, as a random variable normally distributed about retailer-pallet-need with standard deviation ret-pallet-stdDev truncated at 0. Because the method for truncation becomes less efficient the closer the mean is to zero, we display that distance as #stdevs.




Figure 5. Agent behavior dynamics: pallet needs.

The other set of agent behavior dynamics describes the entry and exit of pallet repairers, with parameters shown in Figure 6. The number and spatial distribution of entering firms is controlled by parameters from Figure 4 above; the entry-cycle-type variable in Figure 6 determines whether entry occurs with a regular period of eoc-period-ave or at random intervals. Random intervals can either be uniformly distributed with range [0, 2×eoc-period-ave] or have normal distribution truncated at 0 (and #stdevs is displayed for the same reasons as in Figure 5).




Figure 6. Agent behavior dynamics: repair firms' entry and exit.

The frequency with which repairers consider exit may be identical for every firm, or of two idiosyncratic types, chosen by exit-freq-type: differing medium-runs for each firm or a common medium run length that is shifted a random amount for each firm. The distribution of those idiosyncratic terms is determined by the values of the same parameters that describe exit. This timeframe in Figure 6 is 22 ticks, as in the experiments reported below (see Appendix B).

The decision to leave is based on satisfying an average per-period breakeven quantity (eoc-req-vol). Because taking an average smoothes over the medium run, the results of the experiments that we report below are also affected by the inventory of each repair firm (eoc-capacity). When a firm does exit, eoc-leaving-action specifies whether its inventory of

pallets is sent directly to the landfill or distributed among other repairers during one additional medium-run time period.

There are two groups of pallet dynamics parameters. The first group, in Figure 7, specifies prices and costs and determines pallet classification during processing. Pallets are classified by `new-pallet`, `poor-pallet`, or `unusable-pallet`, based on the number of uses remaining (mimicking typical industry practice). These conditions affect manufacturers' and repairers' decisions on actions to take with their inventory of pallets. Manufacturers replenish their stock at price `new-pallet-price` from an infinite source of pallets and at `used-pallet-price` from repairers. Both firms take into account transportation costs through `cost-mile-empty`, and repair firms also consider per-pallet `repair-cost`. (Firms do not actually calculate profits, they just use these values to decide which actions are more profitable.)




Figure 7. Pallet dynamics: prices, costs, and classification.

The two monitors depicted in Figure 7 display the maximum geographic limits in miles created by the price and cost parameters. From the perspective of manufacturers, “e->m” is the maximum mileage to which manufacturers would go to buy used pallets instead of new pallets: $(\text{new-pallet-price} - \text{used-pallet-price}) / \text{cost-mile-empty}$. From the perspective of repairers, “m->e” is the furthest distance a repair firm will travel to get pallets from the manufacturer: $(\text{used-pallet-price} - \text{repair-cost}) / \text{cost-mile-empty}$. We can abstract away from the party that actually moves the pallet—manufacturer, retailer, or some other contracted entity; what is important is the effect of the tradeoff on agents' decision making.

The second group of pallet dynamics parameters, in Figure 8, specify pallet degradation and repair. Currently there is only one type of pallet, built with an expected number of trips that it can be used before needing repair or sent to the landfill. After each use, a pallet's expected life has a fixed chance (`high-degrade`) of losing two trips instead of one from its useful life. When a repairer takes in used pallets, wood from some of them is used to fix the others, measured by `repair-efficiency`. The repair strategies a repairer may adopt (`eoc-repair-strategy`) mimic practices in the industry: either to repair the best first (leaving the worst pallets to be harvested for parts) or randomly choosing pallets to be harvested for parts.

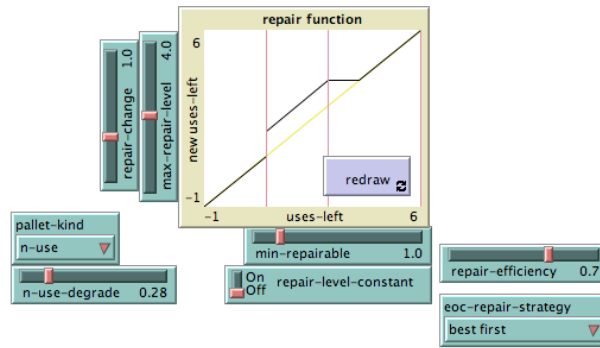


Figure 8. Pallet dynamics: degradation and repair

A graphical display summarizes the parameterization of the pallet repair function, with vertical lines at the poor and unusable pallet cutoffs. The lower and upper limits of possible repair are `min-repairable` and `max-repairable`. The `repair-level-constant` switch chooses whether the repair either brings all pallets to the same (constant) level or causes a fixed increase in the number of uses (the case shown in Figure 8); in the latter case, the fixed increase is given by `repair-change`. The `redraw` button allows the user to dynamically observe the repair function while its parameters are changed.

Agent specification

We now turn to a description of the agents in this inter-firm supply chain: pallets move goods among “users” (that is, from manufacturers to retailers); they are returned either directly to manufacturers or indirectly through End-of-Cycle firms (pallet repairers); and they eventually exit the system in a landfill (another End-of-Cycle firm). Our abstraction of each agent type is influenced by the market’s structure and our interest in investigating open-loop, closed-loop, mixed, and hybrid configurations of the inter-firm supply chain.

Manufacturer agent variables

<code>total-actual-pallet-need</code>	Pallets needed during the current period, determined stochastically during program execution for each retailer link.
<code>inventory</code>	The current set of pallets held by the manufacturer; can be maintained across periods.

Manufacturers are one of the two primary decision makers in the model, ordering new pallets and obtaining used pallets either returned directly from retailers or indirectly through other sources. The number of pallets that a manufacturer requires each iteration is the sum of its retailers’ needs. Each retailer’s pallet need is generated by an independently and identically distributed normal distribution censored at 0, as described above. After receiving a set percentage of pallets back from each retailer (presumably by contract, perhaps an informal one), each manufacturer also receives a scheduled set of new pallets that add to its existing inventory, and if

the resulting inventory does not meet current need, the manufacturer first purchases refurbished pallets from repair firms and then orders additional new pallets to satisfy its need.

Retailer agent variables

inventory	The current set of pallets held by the retailer. Entire inventory expunged each period by returning to manufacturer, sending to end-of-cycle firm, or sending to a landfill.
-----------	--

Retailers generate pallet need as described above (a separate realization of a truncated normal distribution for each link with a manufacturer). The retailers receive their requested number of pallets worth of goods, and those pallets are added to each retailer’s inventory, this agent type’s only variable.

End-of-Cycle third-party repair agent variables

inventory	The current set of pallets held by this third-party repair agent; can be maintained across periods.
pallets-shipped	The total number of pallets sold to goods manufacturers since the agent last considered exit.
active	Indicates that the agent is currently receiving pallets; otherwise, it is getting ready to exit.
period-end	The next time at which this agent will consider exit.
evaluation-period	The time between each of this agent’s exit decisions; only used if each agent has an idiosyncratic medium run length.
market-entry	The time this agent firm entered.
policy-type	This agent’s type of service.

The third-party repair firm has more decision making complexity than either of the previous agents, reflecting this agent’s centrality in determining the model’s outcomes. Third-party firms have inventory that is held across periods, and they track their aggregate pallets shipped to make medium-run exit decisions. To facilitate our tracking of outcomes, these agents also record when they first enter the market. Because repairers exit but may take time to rid themselves of their pallets, they also have a flag indicating whether they are actively accepting pallets or simply disposing of pallets before finally exiting.

This agent currently has a single policy, which is to rebuild all received pallets, removing only the unusable ones from the supply chain. In the future, other policies will be implemented. Including other end-of-cycle types will model hybrids of open- and closed-loop pallet usage. Thus we can model “cooperatives” where pallets are returned jointly by retailers in the same industry (for example liquor stores and local wineries) or the case where the third party actually owns the pallets (CHEP, and smaller “pallet service” providers).

Landfill agent variables

inventory	The set of pallets to be removed from the supply chain this period.
-----------	---

Like retailers, landfills also hold a single inventory variable; at the end of the period their pallet inventory dies, representing either its sale as mulch or its addition to the landfill.

Link agent variables

retailer-links-own	The actual (stochastic) number of pallets needed by the retailer at one end of the link for shipping goods from the manufacturer at the other end of the link.
--------------------	--

The retailer-link breed's single variable is also cleared out each period. It is the current pallet need for shipping from the manufacturer at one end of the link to the retailer at the other end. As explained above, the mean and standard deviation of the retailer-manufacturer links' pallet need are global variables, but the actual need is realized separately for each retailer-manufacturer link.

Pallet agent variables

kind	Pallet type, currently only "n-use".
uses-left	This pallet's remaining number of uses.
repair-cnt	The number of times this pallet has been repaired.
owner	A manufacturer, to whom the pallet will be returned, or empty
curr-location	The user or repair agent where the pallet is currently held.

Finally, the pallet breed has several variables that keep track of the pallet's condition, ownership (if anyone owns it), and location.

Having a separate agent for each pallet is a design element that obviously affects the speed of our simulations, but modeling their deterioration directly (or at least considered as an average of a group of pallets) is appealing. There is an accepted practice of abstracting multiple entities into a single agent, even though that aggregated agent's decisions are designed with a single decision-maker in mind, thus scaling the system down to a reasonable load.

Agent interactions

Before discussing results, we pause to consider two aspects of player interaction in the model. First, the dynamics of relationship between repairers and nearby manufacturers and retailers; second, the paucity of decisions made by retailers.

As described above, manufacturers and retailers may send pallets to repairers if the repairer is close enough to justify that repairer picking up the pallet, and manufacturers may receive pallets from repairers if the cost including transportation is not too high. Each repairer, whether

it is receiving or sending pallets, is able to fill or exhaust its capacity fully before the pallet user moves on to the next-closest repairer. This reflects lower transactions costs to the pallet user when dealing with fewer repairers, as well as lower transportation costs. It is possible for the user in our model to employ the services of several repairers if the user has more pallets than nearby repairers' capacities, or if there is a geographic clustering of users and/or repairers. In the former case the services are all used during the same period. In the latter case the services of different repairers may be employed by a user over different periods because users check with repairers (that are near enough) for available inventory in random order from period to period.

Because pallet users choose from repairers in random order, there can also be mixing of pallet repairers and pallet users from period to period. This is not unrealistic; the actual market does exhibit some of this volatility, particularly in the case of smaller-scale repairers and/or users. In addition, this feature of the model provides some smoothing in the viability of repair firms to compensate for the otherwise drastic drop-off that our model would have on repair firm profitability. In the actual marketplace, pallet users are likely to maintain a relationship with more than one pallet repairer by regularly sending pallets to both of them. Finally, this kind of behavior is more likely to occur in our model when there is a geographic clustering of users and repairers, just as would be true in the real market—where there is little cost difference between pallet repairers, allowing pallet users low-cost diversification.

As we also described above, there is a low degree of sophistication in the retailer agents. In the future, we could endow retailers with decision-making power that reflects an interest in retaining high-quality pallets for sale rather than automatically sending them back to the manufacturer. This would, on the face of it, seem to be the reason for open-loop systems to use lower-quality pallets: manufacturers can't be sure that they will recoup their high-quality pallets, the source of the common resource problem described above.

Yet it is unlikely that retailers are the sources of this market inefficiency: they typically do not want to be distracted by finding another market for their pallets. Instead, when a retailer does not return pallets to their originating manufacturer it may be due to random factors—retailers don't have reason to care to whom each pallet is returned; they don't want to take the time to consider alternative destinations for their used pallets. Their motivation to get rid of the pallets is just to keep them from cluttering up floor space. This would explain existing ownership color coding methods: pallets colored blue (CHEP) or green, pallets with yellow legs (Schneider 2009). These colors make it easier for retailers to identify pallets that should be returned to a certain location, either third parties or manufacturers. The downstream pallet users (in our model, retailers) are contractually obligated to return the pallets, but it appears to be more crucial that it is *easy* for them to return pallets than that they are incentivized to return them.

The source of the different outcomes in open- vs. closed-loop relationships is more likely to be due to the behavior of end-of-cycle firms, the repairers in our model. There is extensive evidence that pallet ownership firms' (CHEP and others) contracts and legal battles are more concerned with the third-party repair firms' incentives than they are with retailers' incentives (Brindley, 2009). We will return to discussing these firms in our conclusions.

As we turn to discussing our experiments, note the asymmetry in the pallet users, the upstream manufacturers and the downstream retailers. The asymmetry has two sources, from the perspective of repairers. The minor source of asymmetry is that retailers have no capacity to

keep pallet inventories across periods, so they are always a source of (non-owned and/or excess manufacturer capacity) pallets for repairers. The major source of asymmetry is that repairers can only sell their pallets to manufacturers.

Experiments

These explorations of parameter space are intended to illustrate how the mixed use of open-loop and closed-loop systems affects the behavior of the system: With regards to the system's dynamics, what regularities are there across parameter settings, what are the differences for manufacturer and retailer pallet users, and what do these results illustrate about our model? There are many parameters in the model, and there are also stochastic components that vary randomly across runs. A full parameter vector specifies a value for all of the sliders and choice menus in Figure 2. As described in the next section, we hold most of those parameters fixed in our explorations below and vary a subset of three parameter values to explore the variations in the system's behavior as that subset of the parameter vector changes. In doing so, we can observe how stylized facts about the market are—or are not—reflected in the model's outcomes. The particular stylized fact that these experiments study is the distribution of pallet repair firms.

Pallet repairers will avoid exit if they aren't too far from a source of pallets and/or not too far from demand for pallets. A manufacturer is likely to attract stable repairers nearby the more pallets the manufacturer owns. Retailers likewise should attract more stable repairers the fewer pallets the manufacturer owns, but there is not a symmetric relationship between the results for manufacturers and retailers. The asymmetry occurs because repairers' income arises only from selling the pallets to manufacturers. It is not sufficient for repairers to take in a large number of pallets from retailers; the repairer also has to be able to sell the pallets to a manufacturer. It's not even necessary for the repairer to take in pallets from a retailer, because manufacturers can be a source of used pallets if they have sufficiently small capacity for holding pallets.

The experiments' parameter values

In these experiments, repair firms enter three (*new-min-eocs*) to four (*new-max-eocs*) at a time, randomly placed in space, and with random frequency in time (*entry-cycle-type*) averaging once every 22 periods (*eoc-period-ave*).

Entry is not contingent on potential profits, a modeling feature not too far from reality because this market requires basic skills held by many individuals in carpentry and construction, and it requires little capital (Freedonia 2008, Schneider 2009). From a modeling perspective, we can also view the entry dynamic as abstractions of both the process that better-informed agents would use to find profitable geographic locations to exploit and the competitive nature of the industry in which established repairers frequently see their business lost to others.

Exit, on the other hand, is contingent, as follows. Rather than model a price-setting market and calculating profit, we use a break-even condition for exit and model the marginal cost curve as horizontal up to a capacity limit where it becomes vertical. If sufficient pallets (*eoc-req-vol*) are not processed during the time since the repairer last considered exit (*eoc-freq-type*,

eoc-period-ave), the repairer exits and disposes of its pallets (eoc-leaving-action). Implementing this cutoff as prices involves simply taking the ratio of fixed costs to variable costs—which, as implied above, can be quite small. Likewise, taxes and subsidies can be incorporated in this cutoff by manipulating the numerator (for lump sum fees or subsidies) or denominator (for *ad valorem* taxes or subsidies) of that fraction.

Appendix B lists the complete set of parameter values for the experiments. There are three that vary across runs: the scale of pallet use in the system (retailer-pallet-need), the balance of the hybrid open-/closed-loop system toward pallet's direct-return to goods manufacturers (%manuf-owned), and repairers' capacity for inventory (eoc-capacity). The last of the three did not affect the dynamics of the system within the range we explored, and so we present results on only the first two below.

Note that every run of the model reseeds before determining the geographic layout, so the fixed distances between goods manufacturers, retailers, and the landfill vary across the repeated runs for a fixed parameter vector: we explore a variety of geographic network topologies.

Results

To gain some context for the results, we present raw time series outcomes before discussing a more compressed visualization of the entire experiment's results. The plots shown in Figure 9 illustrate all the time series generated by one parameter vector run with 32 different random number seed pairs; the individual runs are overlaid, along with bold lines showing the means, first quartiles and third quartiles across the 32 runs for each tick. (The bold lines correspond to the 24th-, 16th-, and 8th-highest values of the 32 time series, at each tick.) The time series in Figure 9 show the evolution over 600 ticks of the average distance from manufacturers to the nearest repairer. One parameter vector's results using 32 different random number seeds is presented in Panel A and another parameter vector in Panel B. The fluctuations in the time series occur as repair firms enter and exit the market; new repair firms may be closer to a manufacturer or retailer than existing repair firms.

For some parameter vectors, as in Panel A of Figure 9, the time series does not settle down, indicating that those parameter values are unreasonable for the model's structure. For other parameter vectors, the distances settle down over time, as in Panel B of Figure 9. Because repair firms are constantly entering at random locations, we can be sure that outcomes like Panel A are distinct from Panel B; the distance from each user to the closest repair firm will be erratic unless there is a reason for one or more to survive near a user. We found 600 ticks to be sufficient time to differentiate between these two types of outcomes using the method in Figure 10. There is still some decrease in the mean value of Figure 9 B's time series between ticks 550 and 600; taking these time series further would heighten the transitions presented in Figure 10.

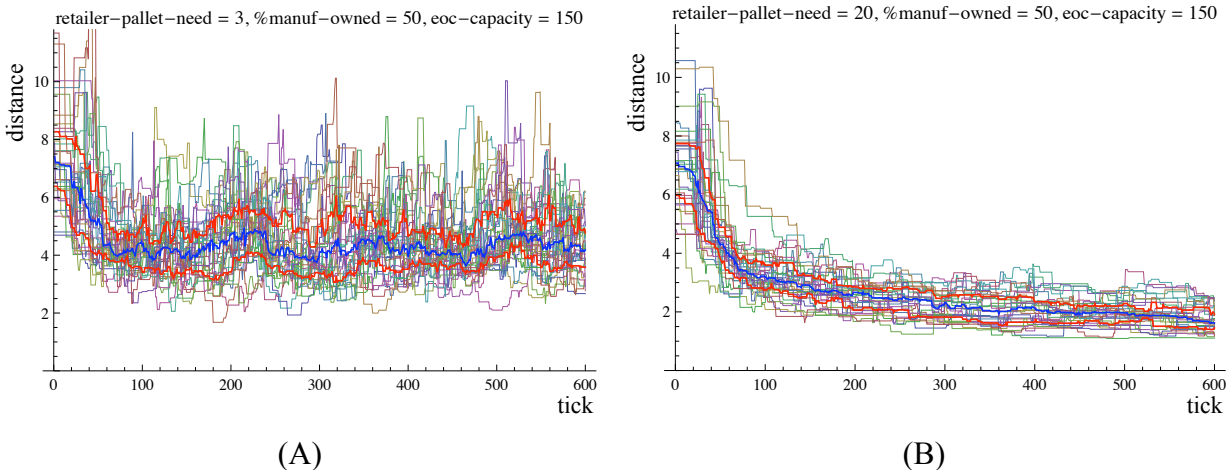


Figure 9. Each panel contains a set of time series generated from 32 runs of the model with all parameters fixed. Panel A shows a case in which repair firms' distance from goods manufacturers does not stabilize. The case shown in Panel B, on the other hand, is one in which the distance from manufacturer agents to the nearest repair firm does not stabilize.

We now present concise summaries of our parameter space explorations. Each plotted point in Figure 10 summarizes the group of 32 runs for each parameter vector; for example, one point in the manufacturer plot of Figure 10 corresponds to the results plotted in Panel A of Figure 9, and another point corresponds to Panel B of Figure 9. The vertical axis in Figure 10 is the average distance from a set of pallet users (manufacturers in Panel A, and retailers in Panel B) to the nearest pallet repair firm, averaged over the final 50 ticks of the 600-tick time series for all 32 runs. Taking the final 50 ticks avoids initialization effects seen in both panels of Figure 9: We are looking for a steady-state response.

Along the horizontal axis of both the manufacturer plot (Figure 10A) and the retailer plot (Figure 10B) is the percentage of pallets that are manufacturer-owned, and the connected points in each plot hold constant the size of retailer pallet needs (varying `retailer-pallet-need` but not `ret-pallet-stdDev`). Plots of the medians are similar; see the notes to Figure 10 for a discussion of standard deviation.

The third variable that we varied, repair firms' inventory capacity (`eoc-capacity`), did not alter the outcomes significantly. Although larger inventory capacity could help a repair firm smooth its sales over the medium run and therefore improve the repairer's profitability, this may not be a factor for the parameter vector in these experiments (Appendix B). The average level of business between manufacturers and retailers (the labels on each plotted line of Figure 10) can also affect repair firms' smoothing of profits across periods, but it affects the profits more directly through the number of pallets available for repair firms in the system each period.

Unstable outcomes, like that pictured in Panel A of Figure 9, are those points in Figure 10 that are toward the top of either plot; they are cases in which the repair firms' locations do not settle down near the pallet user. For example, the outcomes are unstable for all runs in which retailer pallet needs were 3 in both Panels A (manufacturers) and B (retailers) of Figure 10. On the other hand, when retailer pallet need is 4 and manufacturer ownership is 90% or greater, repair firms do settle down to stable locations near manufacturers (Panel A), but they never settle

down to stable locations near retailers (Panel B)—not even for low values of manufacturer ownership.

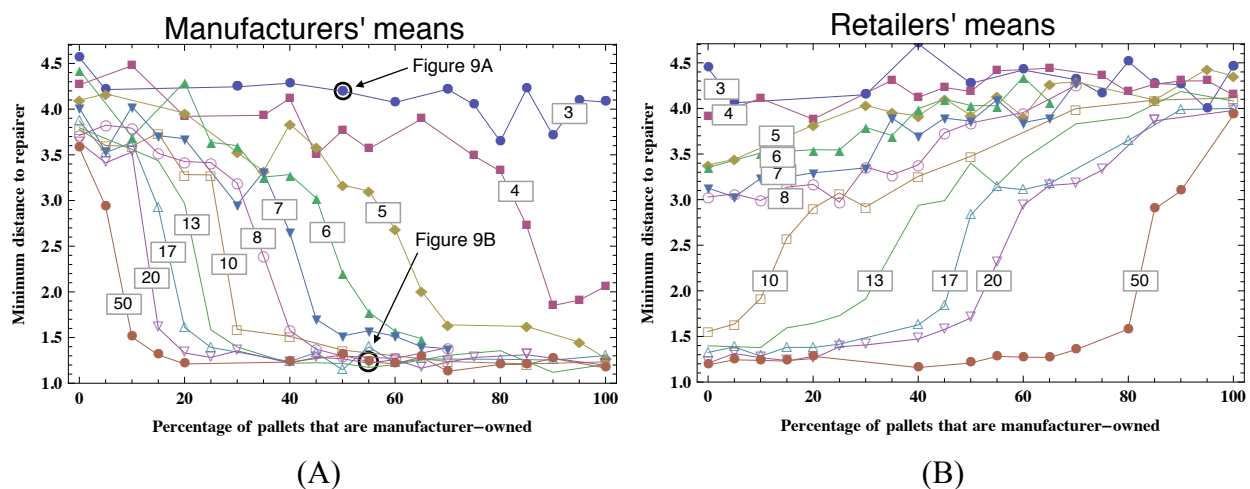


Figure 10. Experiment summaries for (A) manufacturers and (B) retailers show the average distance to the nearest pallet repair firm (vertical axis, measured in pixels). Each plotted point summarizes 32 runs with identical parameters. The two parameters that are varied are the percentage of pallets owned by manufacturers (horizontal axis) and retailer pallet need (boxed labels in the plots). Repair firms are stable near each user type for points in the plot near the horizontal axis. Plotting the standard deviations (not shown) reveals a similar shape: standard deviation is approximately 1 to 1.5 for points higher in these plots (points at about 3.5 or higher); standard deviation falls to near 0.5 or below for the points that are lower in these plots (points at about 1.5 or lower).

For both types of users, we can see that there are more likely to be stable repair firms (y -axis values of 1.5 or lower in the plots) when there is more business between pallet users (when the labels contain higher numbers). For example, when average retailer pallet need is 50, there are stable repair firms near goods manufacturers as long as manufacturers own about 10% or more of the pallets, while repair firms are stable near retailers as long as manufacturers own about 80% or fewer of the pallets. Another example: when retailer pallet need is 20, the range over which repair firms are stable near manufacturers is for ownership of 15% or greater, and near retailers for 50% or fewer.

Thus if the average minimum distance falls for any given level of retailer pallet needs (for any of the connected points) it falls from left to right in the manufacturer plots and from right to left in the retailer plots. In the case of manufacturers, the more pallets that are contractually required to be returned to manufacturers after carrying goods to retailers, the more likely it is that repairers will be geographically tied to manufacturers. The opposite is true for retailers: fewer pallets required to be returned to goods manufacturers means it is more likely that repairers are geographically tied to retailers. Yet the plots show that the retailer and manufacturer graphs are not symmetric.

The asymmetry arises because repair firms find it more difficult to develop stable locations near retailers (few points in the plot have a low mean distance). When the goods transferred between each retailer and manufacturer (therefore the pallets moved between them) is 8 or fewer, stable repair locations cannot develop near retailers—not even with manufacturer ownership as low as 0%. Yet at any level of business above 3 there is a range of manufacturer ownership in

which repair firms find stable locations near the manufacturers. We can explain the asymmetry between by noticing that the repair firms only sell pallets to manufacturers. Yet it is also evident that there are also parameter values for which repair firms do not settle into stable locations near manufactures.

There are four possible conditions for pallet repair firms in the long run: long-term prospects near neither manufacturers nor retailers, near only manufacturers, near only retailers, or near both. These regions are delineated by the two variables depicted in Figure 10, the horizontal axis and the line labels, as they vary in Panels A and B.

- The first region occurs in Figure 10 with pallet need of at most 3 (the labeled line that never drops near the horizontal axis in either panel) for all levels of manufacturer ownership (the horizontal axis). There are other pairs of pallet need and manufacturer ownership that make long run survival impossible: Any points near the top of the plots in both panels of Figure 10. Examples: when pallet need is 4 and manufacturer ownership is 80% or below; when pallet need is 5 and manufacturer ownership is 60% or below; when pallet need is 10 and manufacturer ownership is approximately 15% to 25%.
- Repairers have long-term viability near only manufacturers when pallet need is between 4 and 8 and the minimum manufacturer percentage-owned ranges, respectively, from 90% down to 40%. Other cases in Figure 10 that are close to the horizontal axis in Panel A and far from it in Panel B are also situations in which stable repair firms only exist near manufacturers. They occur for larger percentages of manufacturer ownership (far from the vertical axis) but pallet use that is not too high (smaller label values).
- The region where repairers can survive near both manufacturers and retailers occurs in the opposite case of the second region: when manufacturer ownership is a small percentage but pallet use (the label value) is not too small. This region is much smaller than the others. It contains, for example, pallet use of 10 with manufacturer ownership below 10%; and pallet use of 17 with manufacturer ownership below 20%. In the latter example, more manufacturer ownership would allow repair firms to exist near both retailers and manufacturers, which brings us to the final region of parameter space.
- The fourth region has stable repairers in the long run near both manufacturers and retailers, points that are at 1.5 or less in both Panels A and B of Figure 10. At a pallet use of 50 there is a wide range of manufacturer ownership, 10% to 80%, in which repairers near both pallet user types are viable in the long run. The range decreases as pallet use decreases: the range is only from 20% to 25% when pallet use is 13, and the range collapses at smaller values of pallet use.

The numerical values of these regions are less important than their features. We note that the transitions between these regions, although not immediate (the lines in Figure 10 are not vertical) is fairly quick. The transition is faster (the plot lines' slopes are steeper) for manufacturers. The transition, after averaging out sampling noise, appears monotonic.

The first region, those parameter vectors that make it impossible for pallet firms to settle geographically in the long run, is of most interest. We conclude our discussion of the results with an interpretation of that region of parameter space.

Discussion: validation

We use the repair firms' locations relative to pallet users to explore the viability of parameter settings, a form of model validation. High volatility in repair locations isn't a situation that could last for long in real markets, and thus indicates unrealistic parameter settings. As with any adaptive simulation, we should expect volatility during initial ticks of all model runs; the agents are (or the agent population is) adapting to the environmental parameters. This behavior is in the early periods of both of Figure 9's panels. But as the model run evolves it may or may not approach some measurable steady state. The steady state in this case is some consistency of repair locations, evidencing their ability to maintain sufficient business to break even over the long run.

While exploring the parameter space, we have found a regions where such stability is evident, and these areas represent modeling environments where the forces for change are less likely to modify the functional structure of the model. On the other hand, the region in which repair firms never find a stable distribution would imply long-run changes in the model environment that are not captured by our model's structure. Although one way to view those results would be to recommend modifications to the model's structure, another way to view them is that such changes would miss the point of the exercise. An unstable region is not just erroneous, instead it can tell us something about the features of the model we are exploring, even more than just its limits. While we implemented the method to the extent that Miller (1998) does in exploring a model's fragility, we attempt to capture his spirit in our approach. By retaining the current behaviors, we identify an island of functionality to build upon in future work, and if we were to make modifications it would be only to expand that region of stability.

Conclusions and future directions

These explorations have described our simulation, an abstraction of the shipping pallet inter-firm supply chain. It is an interesting market to study because much of it is organized in an open-loop fashion, because of the common resource problem in these open-loop arrangements, and because most firms mix open- and closed-loop configurations.

The main results of this paper are in its use of a modeling method and validation through stylized facts. The experiments allowed us to characterize the regions of parameter space that do not push the model beyond its useful analog to actual market outcomes. They also described the asymmetries that occur between different pallet user types due to the structure and incentives those agents face.

There are some implications for the pallet industry for geographic regions where the repair market remains fragmented in the manner we model. There are a variety of ownership forms (the horizontal axis of Figure 10's graphs) for which sufficient pallet volume (the vertical axis) can support repeat relationships between pallet users and small pallet repair firms, but the outcome differs for manufacturers (Panel A) and retailers (Panel B). Without sufficient pallet volume, small repair firms have difficulty surviving (outcomes at the top of Figure 10's graphs).

With regards to environmental impacts, there are three sources of functional pallets entering landfills: exiting repairers, manufacturers whose pallets returned by contract temporarily exceed their capacity and local repairers' capacity, and retailers who can't send all their non-contracted pallets to repairers due to temporary lack of inventory space nearby. We will explore these dynamics in future work, and incorporate policy changes (bans, taxes or subsidies, etc.) in our evaluation of the outcomes. Future modeling will also include a second revenue stream for pallet repairers: the option of "chipping" pallets and discarded pallet parts to create garden mulch. It is likely that this mulching is a more important component pallets' end-of-life decisions than landfills are—providing another stylized fact with which we can constrain (validate) parameter choice.

Our model contains agents that are heterogeneous only due to location and idiosyncratic stochastic shocks. We can change a variety of agent decision parameters for all agents, but in future work we will incorporate that variety in decision making at the individual agent level.

We also will incorporate hybrid agents, similar to the "pallet service providers" and "third-party logistics firms" (3PL) that work at small and large scales. When these firms actually own the pallets (CHEP being the canonical example), they have the potential to solve the common resource problem in a fashion familiar to economists—transforming a common resource into a public good. Pallet cooperatives can also mitigate this problem. By incorporating these agents we can then investigate the trade-offs in policies that encourage such behavior or inadvertently discourage it.

Appendix A: Main loop of code

Procedures that take agentset arguments:

```
get-new-pallets, send-pallets-to-eoc, send-pallets-to-landfill
```

Procedure that returns an agentset:

```
worst-overstock
```

to go

```
; Repair firms consider exit
ask eocs with [ticks >= period-end] [ evaluate-eocs-exit-strategy ]

; Repair firms enter
if (ticks mod entry-cycle-period) = 0 [ generate-eocs ]

; Regular delivery
ask manufacturers [ get-new-pallets new-pallet-restock ]

; Current need for pallets
ask retailers [ compute-retail-orders ]
ask manufacturers [
  set total-actual-pallet-need sum [actual-pallet-need] of my-links ]

; Get the necessary additional pallets, first from repairers and then new
ask manufacturers [ get-pallets-from-eoc ]
ask manufacturers [
  get-new-pallets (total-actual-pallet-need - count inventory) ]

; Perform shipping cycle: the empty pallets owned by a manufacturer are returned
ask manufacturers [ ship-goods ]
ask retailers [ return-owner-pallets ]

; Retailers have no capacity for inventory, so send pallets to repairers
ask retailers [ send-pallets-to-eoc inventory ]

; Worst pallets sent to repairers and remainder above total capacity are landfilled
ask manufacturers with [count inventory > manuf-capacity] [
  send-pallets-to-eoc worst-overstock ]
ask retailers [ send-pallets-to-landfill inventory ]
ask manufacturers with [count inventory > manuf-capacity] [
  send-pallets-to-landfill worst-overstock ]

; Repair poor pallets
ask eocs [ eoc-repair-pallets ]

; Landfills dispose of pallets, do plots, and increment the period
landfill-pallets

do-plots
tick
```

end

Appendix B: Parameter settings constant across runs

```
(* Generate network-seed and run-seed later, and use them multiple times *)
set reseed "none (static)"

(* Initial numbers of agents and network characteristics *)
set min-manf 3 set max-manf 3
set min-retailers 10 set max-retailers 10
set distribution-link "all"
set init-min-eocs 10 set init-max-eocs 10
set total-landfills 1
set manf-dist "linear" set retailer-dist "linear" set eocs-dist "linear" set landfill-
dist "linear"

(* Don't actually need next line because linear was chosen above *)
set manf-stdDev 6 set retailer-stdDev 6 set eocs-stdDev 6 set landfill-stdDev 4 set
eocs-stdDev 6

(* Exit and entry *)
set eoc-period-ave 22 set period-stdDev 4.0
set entry-cycle-type "dynamic" set new-min-eocs 3 set new-max-eocs 5
set exit-freq-type "individual phase" set eoc-leaving-action "distribute"
set eoc-req-vol 20" (* set eoc-capacity VARIES *)

(* Shipment-induced drivers *)

(*set retailer-pallet-need VARIES *)
set ret-pallet-stdDev 2 set new-pallet-restock 5

(*set manuf-% owned VARIES*) "set manuf-capacity 30

(* Pallet disposition, movement, and repair *)
set new-pallet 6 set poor-pallet 3 set unusable-pallet 1
set pallet-kind "n-use" set n-use-degrade 0.28

(* Repair function and type *)
set repair-efficiency 0.7
set eoc-repair-strategy "best first"
set max-repair-level 4.0 set repair-change 1.0
set repair-level-constant false set min-repairable 1.0

(* Prices and costs *)
set new-pallet-price 10 set used-pallet-price 3
set repair-cost 0.50 set cost-mile-empty 0.020
```

References

- Axtell, R. L.; Andrews, C. J.; Small, M. J. (2001), "Agent-Based Modeling and Industrial Ecology," *Journal of Industrial Ecology*, 5(4), 10-13
- Brindley, E. C. (2009) "Letter from Ed: Do You Know Where Your Children Are?," *Pallet Enterprise* online feature, <http://www.palletenterprise.com/articledatabase/view.asp?articleID=2986>, posted 12/1/2009.
- Chaturvedi, A. R.; Mehta, S. R.; Dolk, D.; and Ayer, R. (2005) "Agent-based simulation for computational experimentation: Developing an artificial labor market," *European Journal of Operational Research* 166(2005): 694-716.
- Chaturvedi, A. R.; Mehta, S. R.; Dolk, D.; and Gupta, M. (2006) "Computational experimentations in market and supply-chain co-design: a mixed agent approach," *Information Systems and E-Business Management* 4(1): 25-48.
- Choi, T.Y., Dooley, K.J., Rungtusanatham, M. 2001. "Supply networks and complex adaptive systems: Control versus emergence," *Journal of Operations Management* 19: 351–366.
- Coase, R. H. (1960), "The Problem of Social Cost," *Journal of Law and Economics*, 3, 1-44.
- Cooper, T. (2005) "Slower Consumption: Reflections on Product Life Spans and the 'Throwaway Society'," *Journal of Industrial Ecology* 9(1-2): 51-67.
- Fischhoff, B., & Small, M. (2000). "Human Behavior in Industrial Ecology Modeling," *Journal of Industrial Ecology*, 3, 4-7.
- Franke, R. (2009) "Applying the method of simulated moments to estimate a small agent-based asset pricing model" *Journal of Empirical Finance*, 16, 804–815.
- Freedonia Group, Inc. (2008). "Industry Study 2359, Pallets." Project Director: Thomas R. Bowne; Research Analyst: Matt Zielenski. The Freedonia Group, Inc. Cleveland, OH.
- Mabert, V. A. and Venkataramanan, M.A. (1998) "Special research focus on supply chain linkages: challenges for design and management in the 21st century," *Decision Sciences* 29(3): 537-552.
- Gosling, T. (2003) "The simple supply chain model and evolutionary computation," *Proceedings of the Congress on Evolutionary Computation 2003*.

- Gosling, T., Jin, N., and Tsang, E. (2006) "Games, supply chains and automatic strategy discovery", Chapter 18 in J. P. Rennard (ed.), Handbook of Research on Nature Inspired Computing for Economics and Management.
- Hensher, D.A. and Puckett, S. M. (2005) "Refocusing the modeling of freight distribution: Development of an economic-based framework to evaluate supply chain behavior in response to congestion charging," *Transportation* 32: 573–602.
- Koesrindartoto, D., J. Sun, and L. Tesfatsion (2005) "An Agent-Based Computational Laboratory for Testing the Economic Reliability of Wholesale Power Market Designs," *IEEE Proceedings*, Power and Energy Society General Meeting, San Francisco, 931-936.
- Lux, T. (2009) "Stochastic behavioral asset pricing models and the stylized facts" in Handbook of Financial Markets: Dynamics and Evolution, T. Hens and K. Schenk-Hoppé eds., North-Holland: Amsterdam.
- Miller, J.H. (1998) "Active nonlinear tests (ANTs) of complex simulation models," *Management Science* 44(6): 820-830.
- Miller, J.H. and Page, S.E. (2007) Complex Adaptive Systems: An Introduction to Computational Models of Social Life, Princeton University Press: Princeton, NJ.
- Macal, C. M. and North, M. J. (2005) "Validation of an Agent-Based Model of Deregulated Electric Power Markets," presented at the North American Association for Computational and Social Organization (NAACSOS) Conference, Notre Dame, June 26-28.
- North, M. J. and Macal, C. M. (2007) Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation, Oxford University Press: New York, NY.
- Schneider, W. (2009) Vice president for sales, private conversation, 10/2/2009.
- Swaminathan, J.M., Smith, S.F., Sadeh, N.M., 1998. "Modeling supply chain dynamics: A multiagent approach," *Decision Sciences* 29 (3): 607–631.
- Tesfatsion, L. (2007) "Agents come to bits: towards a constructive comprehensive taxonomy of economic entities" *Journal of Economic Behavior and Organization* 63(2): 333-46.
- Thomas, D. J. and Griffin, P.M. "Coordinated supply chain management," *European Journal of Operational Research*, 94(1), 1-15.