3-2010

Local coordination under bounded rationality: Coase meets Simon, finds Hayek

C. Jason WOODARD
Singapore Management University, jason.woodard@olin.edu

Local Coordination Under Bounded Rationality: Coase Meets Simon, Finds Hayek

C. Jason Woodard*
Singapore Management University
jwoodard@smu.edu.sg

5 March 2010

Abstract

This paper explores strategic behavior in a network of firms using an agent-based model. The model exhibits a tension between economic efficiency and the stability of the network in the face of incentives to change its configuration. This tension is to be expected because the conditions of the Coase theorem are violated: the boundedly rational firms in the model lack the ability to discover efficient network configurations or achieve them through collective action. In computational experiments, as predicted by theory, firms frequently became locked into inefficient outcomes or endless cycles of mutual frustration. However, simple institutional innovations such as property rights and side payments dramatically improved outcomes, even for severely myopic firms. The results are consistent with Hayek's observations about the surprising effectiveness of local coordination in the economy at large, but the mechanism involved is different. Instead of prices, firms in the model exploited the fact that their incentives are partially decoupled by the structure of the network. This finding contributes to the growing body of theory on "loosely coupled" interactions in real-world networks.

Keywords: Complex networks, Bounded rationality, Institutional economics, NKC model

JEL Classification: D02, D23, D85, C63, C73

---

*A previous version of this paper circulated under the title "Individual Capabilities and Institutional Conditions in Strategic Network Formation." The paper has benefited greatly from the input of Carliss Baldwin, Jan Rivkin, David Parkes, and Margo Seltzer; the feedback of seminar participants at Harvard Business School, the Academy of Management Annual Meeting, the CCC Doctoral Colloquium, and the Santa Fe Institute Graduate Workshop on Computational Economics; and the comments of an anonymous reviewer. All errors, of course, are my own.

1 Introduction

Although the power and limits of decentralized action have long been central concerns of economics, much is still unknown about the ability of boundedly rational agents to solve complex coordination problems in a socially efficient way. This paper explores the class of coordination problems that arise among firms linked by a supply chain network. In choosing upstream suppliers to buy from (i.e., links to activate), these firms must consider and often anticipate the decisions of their downstream customers. Successful coordination yields efficient and mutually profitable outcomes, while ordering the wrong quantity or combination of inputs leads to missed opportunities or wasted resources. The factors that affect these outcomes are studied using an agent-based model and a set of computational experiments.

In theoretical work, Jackson and Wolinsky (1996) identified "the tension between stability and efficiency" as a central issue for the study of link formation in networks. The tension is twofold. Agents may encounter incentive conflicts that make value-maximizing network configurations unstable, in the sense that they are vulnerable to disruption by agents who form new links or break existing ones. Conversely, inefficient network configurations may be stable, blocking decentralized efforts to achieve more desirable outcomes.

The paper presents a model that exhibits both sides of this tension. On one hand, simulated firms frequently became trapped at locally optimal but globally inferior network configurations. On the other hand, firms were rarely able to settle on global optima when they were encountered. These results are broadly consistent with the Coase theorem, which asserts that agents capable of costless coordination should be able to resolve their incentive conflicts and reach efficient outcomes, but acknowledges that agents subject to transaction costs or other barriers to coordination may not (Coase 1960; Stigler 1966, pp. 113–14).

The firms in the model are unable to fully coordinate their actions because the coordinating mechanisms available to them are limited to adjacent links while the incentive conflicts may span the entire network. Even if they could coordinate more effectively, the firms would not necessarily achieve perfect efficiency because bounded rationality impedes their ability to discover globally optimal network configurations (Simon 1955; Rubinstein 1998).

However, the simulated firms often performed surprisingly well. Firms endowed with better perception or cognition tended to create more efficient networks than less capable firms, and to capture more value in the form of profits. Although these gains were diminished in environments where firms faced more complex incentive conflicts, simple institutional innovations such as property rights (the ability to exclude others from linking to oneself) and side payments (the ability to strike bargains with one's neighbors to obtain mutually agreeable link choices) dramatically improved the observed outcomes, even when the firms' incentives were highly interdependent.

Counterintuitively, these institutions yielded more efficient networks not by improving the firms' collective ability to search for high-value configurations, but simply by facilitating local coordination among customers and suppliers. The ability to influence the decisions of their immediate neighbors often allowed firms to avoid endless cycles of mutual frustration that could otherwise propagate through the network. While most of these firms were only partially satisfied with the resulting network configurations, the value of achieving a stable though suboptimal configuration tended to exceed that of jumping from one frustrated state to another in search of more desirable alternatives.

This finding recalls Hayek's observation about the price system: "The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action" (1945, pp. 526–27). But here there are no prices; instead firms exploited the fact that their incentives are partially decoupled by the structure of the network. While this is a strong assumption, it faithfully models what Weick (1976) called "loosely coupled systems," which arise commonly in social and economic settings.

The rest of the paper is structured as follows. Section 2 presents the model and discusses its relationship to previous work. Section 3 reports results for the basic version of the model and explains the observed alignment between efficiency and stability. Section 4 reports the main results for the extended model, which show that stronger individual capabilities and richer institutional conditions lead to better economic performance at the firm and industry level. Section 5 interprets the results and elaborates on their robustness and limitations. The paper concludes with a summary and brief discussion of opportunities for future research.

2 Model

Consider a fixed set of n firms, denoted $N = \{1, \ldots, n\}$. The firms are arranged in a network whose links represent possibilities for repeated economic interaction. For concreteness, these interactions are described in terms of supply-chain relationships: nodes represent customers and suppliers, and links represent purchase orders.

The basic version of the model is a member of the *NKC* family of fitness landscape models (Kauffman and Johnsen 1991; McKelvey 1999), which were originally used to study evolutionary dynamics in biological ecosystems. Here these dynamics take place in the context of a business ecosystem (cf. Iansiti and Levien 2004) where a firm's strategic choices may reshape the incentives of the firms it interacts with. Navigating these "dancing landscapes" (Kauffman 1993, p. 243) is typically more challenging than optimizing in a static environment because a firm's optimal choice may be contingent on the choices of other firms.

After presenting the basic model, two kinds of extensions are introduced. First, firms are allowed to vary in their individual capabilities, specifically their ability to perceive the same-period linking choices of other firms and to evaluate alternative combinations of links. Second, two institutional innovations enable firms to improve their ability to coordinate their actions. Both versions of the model were implemented using agent-based modeling software and instantiated in a series of computational experiments. The experiment design is summarized at the end of this section.

2.1 Basic model: Iterated link choices with random payoffs

As in other *NK*-style models (e.g., Levinthal 1997; Rivkin 2000; Press 2007), firms interact on a fitness landscape characterized by a set of evolving elements, an associated state space, and a mapping from states to numerical fitness values, which in turn influence the evolution of the elements. Here the elements are the network links among the firms, each of which is either active or dormant in a given period. Activating a link represents a decision by one firm to transact with another, for example by submitting a purchase order. A firm's fitness is determined by its upstream link choices (i.e., which suppliers it buys from) conditional on the choices of its downstream neighbors (i.e., which customers buy from it). The evolution of the network is driven by the firms' repeated application of a best-response decision rule.

Network topology with tunable complexity

As in the real world, the structure of the network may be constrained by restrictions on link formation. For example, transportation costs may rule out transactions between firms that are physically too far away from each other. Or firms may produce goods that are used in different stages of a production process, so that only those immediately upstream and downstream of each other need or want to transact. To model these constraints, each firm can only link to a subset of the others. For each firm $i \in N$, let $U_i \subseteq N$ denote the set of firms to which i can link (its upstream suppliers). Let $D_i = \{j \in N : i \in U_j\}$ denote the set of firms that can link to i (its downstream customers). The *topology* of the network is defined by $\Gamma = (U_1, \ldots, U_n)$.

Figure 1 shows an example of a network topology in which each firm has exactly two

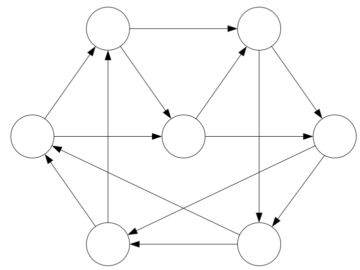

Figure 1: A network topology with two potential customers and suppliers per firm.

customers and two suppliers. In the experiments, the topology for each simulation trial is generated by forming a regular lattice: arranging the n firms in a circle and connecting each firm to its clockwise neighbor, then its neighbor's neighbor, and so on until each firm has exactly k downstream and k upstream neighbors.[1]

Iterated sequence of interactions among firms

The firms interact sequentially over a set of periods $T = \{1, \ldots, \tau\}$. In each period $t \in T$, each firm observes the state of the network and chooses a set of links to activate. The decisions of all firms, taken together, determine the state of the network at the end of the period.

Each firm may choose to order from (i.e., activate a link to) any subset of its upstream suppliers. Fulfilling the order may be costly for one or both parties, for example if the supplier needs to customize its product or the customer needs to adapt its production process for the supplier. A customer may choose to renew the order in a subsequent period, but for simplicity this situation is treated the same as if the order were placed for the first time, with no additional costs or benefits. The actions available to firm i are thus given by the power set of U_i, denoted $\mathcal{U}_i = 2^{U_i}$. Similarly, the downstream actions that affect i's payoff are given by $\mathcal{D}_i = 2^{D_i}$. Let $U_i^t \in \mathcal{U}_i$ be the set of firms to which i links in period t, and $D_i^t = \{j \in N : i \in U_j^t\}$ be the set of firms that link to i in that period. All links are initially dormant, i.e., $U_i^0 = D_i^0 = \varnothing$ for all $i \in N$.

Payoff tables with random interactions among links

At the end of each period, each firm receives a numerical payoff representing the economic value it derives from transacting with its upstream and downstream neighbors during that period. The payoff to firm i in period t is denoted π_i^t. While in principle a firm's payoff might depend on the link choices of any other firm (or factors unrelated to the state of the

[1] Preliminary experiments studied networks with random topologies; the results were qualitatively identical.

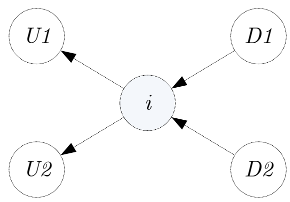

$D1$	$D2$	$U1$	$U2$	π_i
0	0	0	0	.37
0	0	0	1	.26
0	0	1	0	.59
0	0	1	1	**.92**

$D1$	$D2$	$U1$	$U2$	π_i
0	1	0	0	.11
0	1	0	1	**.83**
0	1	1	0	.50
0	1	1	1	.64

$D1$	$D2$	$U1$	$U2$	π_i
1	0	0	0	.49
1	0	0	1	**.66**
1	0	1	0	.20
1	0	1	1	.38

$D1$	$D2$	$U1$	$U2$	π_i
1	1	0	0	.02
1	1	0	1	.41
1	1	1	0	**.87**
1	1	1	1	.75

Figure 2: A payoff table for a representative firm with two customers and two suppliers.

network), here we assume that it depends only on the set of suppliers the firm chooses to transact with and the set of customers that choose to transact with the firm. In other words, a firm's payoff is independent of the choices of firms that are not its customers, and also of its customers' choices with respect to their other suppliers. Formally, let i's payoff at time t be given by $\pi_i^t = \pi_i(U_i^t; D_i^t)$. The function $\pi_i : \mathcal{D}_i \times \mathcal{U}_i \rightarrow \mathbb{R}$ is called a *payoff table*. The profile of payoff tables $\pi = (\pi_1, \ldots, \pi_n)$ defines the *value function* for a particular instance of the model (cf. Jackson 2004).

Figure 2 shows an example of a payoff table for one of the firms in the previous figure. The full table has 16 entries; it is divided into four sub-tables, one for each possible configuration of downstream links. The relative magnitudes of the payoffs reflect the pattern of complementarities among the links. Two links are complements (substitutes) if the presence of one increases (decreases) the marginal value of the other, where the marginal value is the difference between the payoff with and without the link activated. The payoff associated with the best response for each configuration of downstream links—that is, the best choice of suppliers given the customers that have ordered in the given period—is shown in bold type.

In the experiments, the payoff table entries for each simulation trial are determined by independent draws from a uniform distribution on the unit interval. This procedure expresses the assumption that each link in which a firm participates is equally likely to be a complement or substitute for every other. According to Kauffman (1993, p. 41), this is equivalent to "confess[ing] our total ignorance and admit[ting] that … essentially arbitrary interactions are possible" among the links.

A myopic best-response decision rule

The tuple (N, Γ, π) defines a game in strategic form.[2] While the pure-strategy Nash equilibria of this game will play a role in explaining the observed dynamics of network formation, we are interested in modeling firm behavior that may or may not lead to a game-theoretic equilibrium. We therefore need to define the firms' decision rule explicitly.

Both the basic and extended versions of the model use a myopic best-response rule: given the most recently observed actions of its customers, each firm evaluates its own potential actions within the limits of its cognitive abilities, then selects the action that would yield the highest next-period payoff if its customers' actions remain unchanged. In the basic model, the most recently observed actions are those of the previous period, and firms are capable of considering all 2^k supplier combinations in evaluating their expected payoffs. Firm i's action is given by $U_i^t = \arg\max_{U \in \mathcal{U}_i} \pi_i(U; D_i^{t-1})$, recalling that $\mathcal{U}_i$ is the power set of suppliers and D_i^{t-1} is the set of customers that ordered from i in the previous period.

2.2 Extended model: Individual capabilities and institutional innovations

The basic model assumes, in effect, that firms act simultaneously in each period and are capable of considering an unlimited number of supplier combinations. The extended model relaxes these assumptions by parameterizing two important firm capabilities, labeled perception and cognition. By varying these parameters experimentally, we can simulate firms that differ in their ability to collect and process information.

Even with these additional parameters, firms cannot interact with each other except through their link choices. Their ability to coordinate is thus severely limited. For example, a supplier may be harmed if it accepts an order from a particular customer; perhaps the total order volume would exceed the supplier's capacity, or the customer poses a credit risk. In the basic model, the supplier would have no choice but to accept the harm. To study the effects of alternative institutional conditions that support richer modes of interaction, the extended model introduces two institutional innovations, property rights and side payments.

Individual capabilities

The perception parameter, ρ, allows firms to observe a fraction of their peers' current-period actions in addition to the previous-period actions of all firms. The cognition parameter,

[2] A strategic game consists of a set of players, a set of actions for each player, and a set of payoffs over profiles of actions (Osborne and Rubinstein 1994, pp. 11–13). Recall that the actions $\mathcal{U}_i$ are simply the power sets of U_i, which are defined by Γ.

γ, restricts the number of alternatives that a firm can evaluate in a given period. These parameters can either be fixed for all firms in a simulation trial or varied individually by firm (substituting ρ_i for ρ and γ_i for γ in the definitions below).

Perception In the extended model, what a firm sees when it is called to choose an action is determined by its *frame*, a concept inspired by Goffman (1974). Firm i's frame, denoted F_i^t, is a set containing the firms whose previous actions in period t are visible to i. The perception parameter determines the contents of a firm's frame as follows. Let $F_i^t = \{j \in N : j < \min(i, \rho n)\}$, where perception $\rho \in [0, 1]$. Then let $\widehat{D}_i^t$ be the set of downstream neighbors perceived by i to have activated links with it when i is called to take a decision during period t, i.e., $\widehat{D}_i^t = \left\{(D_i^t \cap F_i^t) \cup (D_i^{t-1} \setminus F_i^t)\right\}$. Note that setting $\rho = 0$ is equivalent to simultaneous interaction (as in the basic model), i.e., $F_i^t = \varnothing$ for all i. At the other extreme, $\rho = 1$ yields sequential interaction with perfect information, i.e., $F_i^t = \{j \in N : j < i\}$.[3]

Cognition Recall that the number of actions in each firm's choice set grows exponentially with k, since the best-response rule requires a firm to evaluate 2^k rows of its payoff table to determine the action that maximizes its payoff given the perceived network state. The cognition parameter captures the idea that a firm may only be able to evaluate a subset of its potential actions in each period. For $\gamma \in [0, 1]$ a firm randomly selects a set $\widehat{\mathcal{U}}_i^t \subseteq \mathcal{U}_i$ of approximately $\gamma 2^k$ actions to evaluate in period t. This set always includes the "status quo" action, U_i^{t-1}. Note that $\gamma = 1$ corresponds to the basic model, where $\widehat{\mathcal{U}}_i^t = \mathcal{U}_i$. For $\gamma = 0$, firms are trapped in their initial configuration, i.e., $\widehat{\mathcal{U}}_i^t = \left\{U_i^{t-1}\right\} = \{\varnothing\}$ for all $t \in T$.[4]

Taking into account the downstream link states perceived by the firm, $\widehat{D}_i^t$, and the set of upstream actions available to evaluate, $\widehat{\mathcal{U}}_i^t$, the best-response decision rule for the extended model (without property rights or side payments) is $U_i^t = \arg\max_{U \in \widehat{\mathcal{U}}_i^t} \pi_i(U; \widehat{D}_i^t)$.

[3] To reduce the possibility of simulation artifacts, the order in which the firms act is randomized at the beginning of each period (Axtell 2000). Formally, this entails renumbering the firms using a permutation function. This function is omitted from the text to reduce notational clutter.

[4] I also ran experiments in which $\gamma = 0$ yielded random behavior. This yielded similar efficiency outcomes but starkly different dynamics because the firms wandered at random over the landscape instead of remaining at their initial configuration. Due to subtle interactions between this behavior and the institutional innovations, I decided in favor of the simpler rule.

Institutional innovations

An institution is a social technology in the sense of Nelson and Sampat (2001): a pattern of productive activity analogous to a physical technology but involving human interaction rather than physical engineering. Under the institutional conditions of the basic model, which I label HOBBESIAN to suggest a stylized "state of nature," firms lack the ability to block the actions of their downstream customers or compensate their upstream suppliers in exchange for making or breaking a link. Although these conditions may seem implausible in a modern market economy, they are consistent with the existence of long-term agreements between customers and suppliers that require suppliers to accept orders placed at any time, as well as laws or regulations that require suppliers such as utility companies to serve all customers, even if some are unprofitable.

The extended model explores two institutional innovations that more closely resemble those available to a typical present-day firm: the right to refuse orders from customers, labeled PROPERTYRIGHTS, and the ability to transfer a numeraire good (money) to suppliers in exchange for accepting an order, labeled SIDEPAYMENTS. Like the perception and cognition parameters, these institutions alter the firms' decision rule but leave the other aspects of the model unchanged.

Property rights The PROPERTYRIGHTS institution grants upstream firms the right to refuse links using the following protocol. When evaluating the payoff associated with a particular configuration of links (i.e., a row in its payoff table), a downstream firm asks each supplier to whom it is proposing a link to consider the proposed configuration. Each of these suppliers returns a "yes" if its anticipated best-response payoff for the proposed configuration equals or exceeds its anticipated payoff for the status quo (consistent with its level of perception and cognition), or a "no" otherwise. If all firms approve, the proposed configuration remains in the set of actions under consideration, otherwise it is discarded.

Let $\widetilde{\mathcal{U}}_i^t \subseteq \widehat{\mathcal{U}}_i^t$ be the set of actions that obtain unanimous approval in period t. The best-response rule for the PROPERTYRIGHTS case is then $U_i^t = \arg\max_{U \in \widetilde{\mathcal{U}}_i^t} \pi_i(U; \widehat{D}_i^t)$.

Side payments The SIDEPAYMENTS institution grants downstream firms the ability to bargain with their suppliers, potentially inducing them to accept links they would otherwise refuse. It is modeled as a direct extension of PROPERTYRIGHTS. Rather than asking for a yes or no answer when presented with a link proposal, firm i asks each supplier j for the marginal value of its proposed action—that is, the difference between its anticipated best-response payoff with and without a link from i to j. Again, we assume that the suppliers respond truthfully within the limits of their perception and cognition.

Let $\widetilde{U}_i \in \widehat{\mathcal{U}}_i^t$ be an action proposed by firm i. For each supplier $j \in \widetilde{U}_i$, let Δ_j denote j's perceived marginal value for the link in question (i.e., the difference between j's anticipated best-response payoff and its anticipated payoff for the status quo), and similarly for Δ_i. Then $\Delta = \Delta_i + \sum_j \Delta_j$ is the total perceived gains from trade for the proposed action. If this action is taken, the firms split these gains equally, according to the Nash bargaining solution (Nash 1953). This is achieved by firm i making a transfer payment of $\Delta/(\left|\widetilde{U}_i\right| + 1) - \Delta_j$ to each supplier j. (The transfer payment will be negative if the supplier expects to gain from the proposed action.) The modified best-response rule for the SidePayments case selects the action with the highest anticipated gains from trade, i.e., $U_i^t = \arg\max_{U \in \widehat{\mathcal{U}}_i^t} \Delta(U; \widehat{D}_i^t)$.

It is reasonable to imagine that real-world firms might respond strategically to link proposals rather than revealing their true valuations. The assumption of truthful behavior can thus been seen as a best-case scenario. However, prior work has shown that the Nash bargaining solution is robust to a wide range of assumptions about the details of the bargaining process (Binmore et al. 1986), albeit subject to more stringent assumptions on the firms' knowledge and rationality. It therefore seems like an appropriate point of departure, which can undoubtedly be improved upon by future work. Fortunately the main results do not depend on the exact amounts of the transfer payments, as long as the realized outcome is the one that is perceived as efficient by the focal firm and its suppliers.

2.3 Experiment design

I implemented the model in Java using a hybrid of the Repast and MASON agent-based modeling toolkits (North et al. 2006; Luke et al. 2004), and conducted three computational experiments. The first explored the dynamics of the basic model in detail, with a focus on the efficiency and stability of the networks formed by the firms. The second studied the extended model under the assumption of homogeneous firm capabilities. The third relaxed the homogeneity assumption to study the behavior and performance of heterogeneous firms. Table 1 summarizes the experiments and simulation parameters.

Simulation parameters

Each experiment consisted of a series of trials. The length of each trial, τ, was 1,500 periods to ensure convergence of the network dynamics. (Only in rare cases were transient dynamics

Experiment / Model	Number of firms (n) Suppliers per firm (k)	Perception (ρ) Cognition (γ)	Institutional conditions	Trials per parameter combination	Total number of observations
1 Basic	{10, 20, 30} {1, 2, 3, 4}	0 1	Hobbesian	1,000	12,000
2 Extended	{10, 20, 30} {1, 2, 3, 4}	{0, 0.25, 0.5, 0.75, 1} {0, 0.25, 0.5, 0.75, 1}	{Hobbesian, PropertyRights, SidePayments}	100	90,000
3 Extended	20 3	Uniform[0,1] Uniform[0,1]	{Hobbesian, PropertyRights, SidePayments}	10,000	30,000 trials 600,000 firms

Table 1: Computational experiments and simulation parameters

longer than 1,000 periods encountered.) The trials varied the number of firms, n, and the number of inputs and outputs per firm, k, as well as the firms' individual capabilities (perception, ρ, and cognition, γ) and institutional conditions (Hobbesian, PropertyRights, and SidePayments). In the first and second experiments, all firms in a given trial were endowed with the same capabilities. In the third experiment, perception and cognition were drawn uniformly at random for each firm.

Each trial was run with a different random seed, yielding a total of 132,000 different model instances (distinct combinations of network topology and payoff tables). I chose values of n large enough to make the firms' decision problem analytically intractable but small enough to find the globally optimal network state computationally, as described below. The choice of k was similarly pragmatic, since optimal solutions were frequently difficult to find for $k > 4$. Consistent with Kauffman (1993), however, I found no qualitative differences in model behavior for k values between 3 and 5, and no evidence to suggest that higher values of k would have produced results that contradict the main findings of the paper.

Outcome variables

For each trial, three outcome variables were recorded: the average network value, average economic efficiency, and average outcome stability achieved by the firms. These within-trial variables are defined below; the results sections discuss their distribution across trials.

Network value The most natural measure of the firms' collective performance is the average value of the networks they form. Let $v(t)$ denote the total network value at the end of period t of a given trial:

$$v(t) = \sum_{i \in N} \pi_i^t.$$

For comparison across trials, reported values are normalized by the number of firms. Let $\textsc{Val}(t)$ be the normalized network value for period t, then define the average network value for a trial as the mean of $\textsc{Val}(t)$ over a set of periods $\mathcal{T}$:

$$\begin{aligned} \textsc{Val}(t) &= \frac{v(t)}{n} \\ \textsc{AvgVal} &= \frac{1}{|\mathcal{T}|}\sum_{t\in\mathcal{T}} \textsc{Val}(t). \end{aligned}$$

In reporting the results, time averages are computed over the last third of each trial ($\mathcal{T} = \{1001, \ldots, 1500\}$) to reflect as much as possible the steady-state behavior of the system.

Economic efficiency Both $\textsc{Val}(t)$ and $\textsc{AvgVal}$ range over the unit interval, like the individual firm payoffs. However, a value of one is unachievable unless each firm is fortunate enough to have drawn a payoff table entry of exactly 1.0, all of which are attained simultaneously—an event with effectively zero probability. Since the payoffs are chosen randomly for each instance of the model, we need a measure of efficiency that relates the realized value in a given period to the maximum *feasible* value for that instance, denoted v^*.

We could naively define v^* by adding up the highest entries in each firm's payoff table, i.e., $\sum_{i\in N} \max_{D\in\mathcal{D}_i, U\in\mathcal{U}_i} \pi_i(U; D)$. But this value is also unlikely to be feasible. Consider two neighboring firms i and j. It may be the case that i's highest payoff is for a state that includes a link to j, but j's highest payoff requires the absence of a link from i. Thus, there are many combinations of individual payoffs that do not correspond to feasible networks.[5] Computing the optimal value therefore requires solving the combinatorial optimization problem

$$v^* = \max_{G\in\mathcal{G}} \sum_{i\in N} \pi_i(U_i(G); D_i(G)),$$

where $G = (U_1, \ldots, U_n)$ is a feasible network state, $\mathcal{G} = \mathcal{U}_1 \times \ldots \times \mathcal{U}_n$ is the set of all feasible states, and $D_i(G)$ and $U_i(G)$ pick out the downstream and upstream links associated with firm i in state G. This is a computationally hard problem. In the experiments, it is solved for each model instance using the commercial optimization package CPLEX, which unlike the firms in the model is able to compute a globally optimal state for each model instance.

Using v^*, it is straightforward to define the efficiency of a network at a given period, and the average efficiency for a trial:

[5] There are $\left(2^{2k}\right)^n = \left(2^{nk}\right)^2$ ways to add up the firms' payoffs, but only 2^{nk} distinct network states.

$$\begin{aligned} \textsc{Eff}(t) &= \frac{v(t)}{v^*} \\ \textsc{AvgEff} &= \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \textsc{Eff}(t). \end{aligned}$$

Outcome stability Economic theory alerts us to the possibility of tension between efficient and stable network states (Jackson and Wolinsky 1996): efficient states may be unstable and vice versa. To explore this idea in a computational setting, we need an observable measure of stability that corresponds to the intuitive concept of "churn" in the network dynamics.

The Nash equilibrium offers a natural starting point. In the basic model, if the same state is realized in two consecutive periods, we can infer that the actions of the firms are a pure-strategy Nash equilibrium of the strategic game induced by the model. This is because in the basic model each firm plays a best response to the state of the network in the previous period. If these actions leave the state unchanged, it is because no firm preferred an action other than the one it took in the previous period; then by definition the actions are a (strict) Nash equilibrium. Moreover, because the best-response rule is deterministic, Nash equilibria correspond to absorbing states of the network dynamics: once reached, an equilibrium state will remain unchanged for all subsequent periods.

Anticipating the results, we will see that stability defined as a Nash equilibrium is strongly related to efficiency. However, two considerations motivate a more general measure. First, it will not always be the case that absorbing states correspond to Nash equilibria of the simultaneous game. For example, with the PropertyRights rule in effect, the firms could get stuck at a state where at least one firm would prefer to choose a different action, but no such action is accepted by the upstream neighbor(s) whose permission is required. Second, even if the network never completely settles down to a single state, parts of it may remain "frozen" from period to period. We would therefore like a measure that works consistently across the basic and extended models and accounts for partial stability at the firm level.[6]

The proposed measure, which I call *outcome stability*, is the fraction of firms in a given period that are "satisfied" in the limited sense that their expectations are fulfilled. Recall

[6]Kauffman (1993, p. 203) explores the emergence of "frozen cores" in Boolean networks, which are closely related to the basic model. My stability measure differs from Kauffman's in that it considers all agents whose decisions remain the same over time, regardless of whether they form a connected component of the network. Separately, Rivkin and Siggelkow (2002) provide a helpful discussion of "sticking points" and their relationship to Nash equilibria. In both their model and the present one, what drives a wedge between a sticking point and a Nash equilibrium is the assumption of an unmodeled subgame—in their case a manager's decision to submit alternatives to a CEO, and in ours an upstream neighbor's decision to accept or reject a link.

that firm i bases its action in period t on the expectation that the set of active downstream links it perceives when called to take a decision, $\widehat{D}_i^t$, will be the same as those in effect when payoffs are awarded at the end of the period, D_i^t. (Also recall that in the basic model, $\widehat{D}_i^t = D_i^{t-1}$; all firms perceive the state of the network at the end of the previous period.) This condition is captured by defining a function $s_i(t) = \Big\{1 \text{ if } \widehat{D}_i^t = D_i^t, \text{ 0 otherwise}\Big\}$. Then let outcome stability and its per-run average be defined as follows:

$$\begin{aligned} \mathrm{OS}(t) &= \frac{1}{n}\sum_{i\in N} s_i(t) \\ \textsc{AvgOS} &= \frac{1}{|\mathcal{T}|}\sum_{t\in\mathcal{T}} \mathrm{OS}(t). \end{aligned}$$

By these definitions, outcome stability will be one in a given period if all firms in that period are satisfied—that is, if they receive the payoffs they expected when they chose those actions using their best-response rules. In the basic model, this is exactly the condition for a Nash equilibrium of the stage game. At the opposite extreme, an outcome stability of zero implies that every firm was "frustrated" by the choices of the others.

3 Results: Efficiency and stability

This section examines the results of Experiment 1. The goals of this experiment were (a) to identify the factors associated with the formation of high-value networks, and (b) to explain the causal mechanisms by which these factors operate in the basic model. As one would expect based on similar NK-style models, I found that the number of suppliers per firm, k, plays a crucial role. In simulation trials with higher k, the interdependent choices of customers and suppliers came into conflict more frequently, yielding networks that were both lower in absolute value and less efficient relative to the optimal value than trials in which firms interacted with fewer network neighbors. In other words, the strategic complexity of the agents' coordination problem had a negative effect on their collective performance.

While this result is intuitive and straightforward, the causal mechanisms involved are more subtle. Specifically, high-value networks tended to occur in trials whose dynamics converged to equilibrium (or, more generally, attained higher outcome stability). Although stable networks need not be efficient in general, they are on average in this model due to the statistical properties of the random payoffs. Section 5 discusses the extent to which these properties are realistic. The focus of this section is to develop a clear understanding their role in the basic model, to build confidence in the analysis of the extended model that follows.

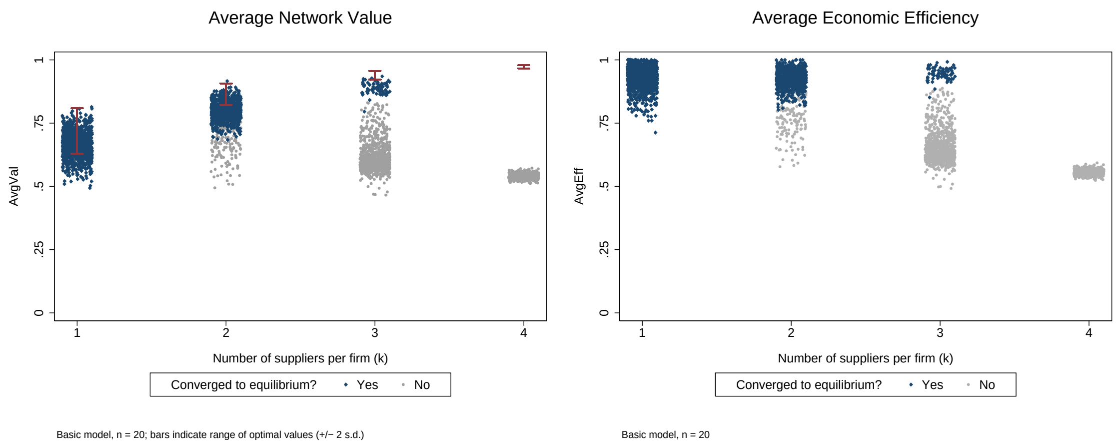

Figure 3: Average network value and economic efficiency (fraction of optimal network value) as a function of the number of suppliers per firm (k). Trials that converged to equilibrium are highlighted.

3.1 From network value to efficiency

Figure 3 summarizes the relationship between the number of upstream links available to each firm (k), and the average network value and efficiency (AvgVal and AvgEff, as defined above). Each point represents a single trial for the case of $n = 20$. The vertical bars in the left-hand graph indicate the distribution of optimal values (v^*) for the trials shown. Trials that converged to equilibrium (i.e., in which AvgOS $= 1$) are shaded in dark color, while those that did not converge are shown in light gray.

Three features of these graphs are immediately apparent:

- The average optimal value is monotonically increasing in k, and converging toward one. Intuitively this is because networks with higher k can take on a larger number of states, yielding more combinations of feasible configurations and thus a higher expected maximum among them.

- The average value of the actual networks formed is concave in k, with a maximum at $k = 2$. This "inverted U" shape is a consequence of the interaction between the value of the optimal configuration, which tends to increase with k, and the difficulty of sustaining this configuration (or a similarly valued one), which also increases due to the higher incidence of incentive conflicts among the firms.[7]

[7]The fact that the maximum occurs at $k = 2$ is not easily predictable from theory, though it may be significant that prior work on random Boolean networks has found this to be a "critical value" of complexity (Derrida and Pomeau 1986), and that between values of 2 and 3 there exists a phase transition between regimes of ordered and chaotic network dynamics (Gershenson 2004).

- Average efficiency is monotonically decreasing. This indicates that when the effects of k on optimal and average network values are combined, the result is unambiguous: even where an increase in k allows firms to create networks with higher absolute value (i.e., in shifting from $k = 1$ to $k = 2$), on average these networks realize a lower fraction of the optimal value (the sum of the payoffs that the firms would receive if a benevolent dictator were to impose the optimal configuration of links).

The results were nearly identical for $n = 30$. For $n = 10$ there were a small number of trials with higher value and efficiency for $k = 4$. This suggests that sustaining high-value network configurations in settings of greater strategic complexity becomes more difficult when there are more firms in the network.

If our analysis of the basic model were to end here, we would have characterized the response of the key outcome variables to the experimental parameters, but we would fail to understand *why* higher values of k tend to yield less efficient networks. The explanation is revealed by the shading of the trials: higher-value networks are clearly associated with trials that converged to equilibrium, and these trials become rarer as k increases. This motivates us to turn our attention to the third observed variable, outcome stability.

3.2 The alignment between efficiency and stability

The left-hand graph of Figure 4 shows a plot of average outcome stability versus network value. Recall that convergence to equilibrium corresponds to AvgOS = 1, so the points lined up on the right side of the graph correspond to the trials that converged in Figure 3. This picture reveals a deeper relationship between stability and collective performance: for trials with the same level of strategic complexity, the average network value is highly correlated with outcome stability, with a slope that increases with k.

This pattern admits a simple explanation. In an equilibrium (AvgOS = 1), all firms are satisfied in the sense that for every period after the cutoff, the actions they choose in response to the previous period's network state are still viewed as best responses *ex post*, when payoffs are awarded at the end of the period. The average value of an ex post best-response action is easy to calculate: it is the expectation of the highest of 2^k independent draws from a standard uniform distribution, or $\frac{2^k}{2^k+1}$. On the other hand, if a chosen action is not an ex post best response, let us suppose that the associated payoff could have come from any other row in the payoff table, for an average value of 0.5, the ensemble average over all payoffs.[8]

[8] In fact, the true average will be slightly lower than the ensemble average because, by assumption, the payoff received is *not* the one that was the ex ante best response.

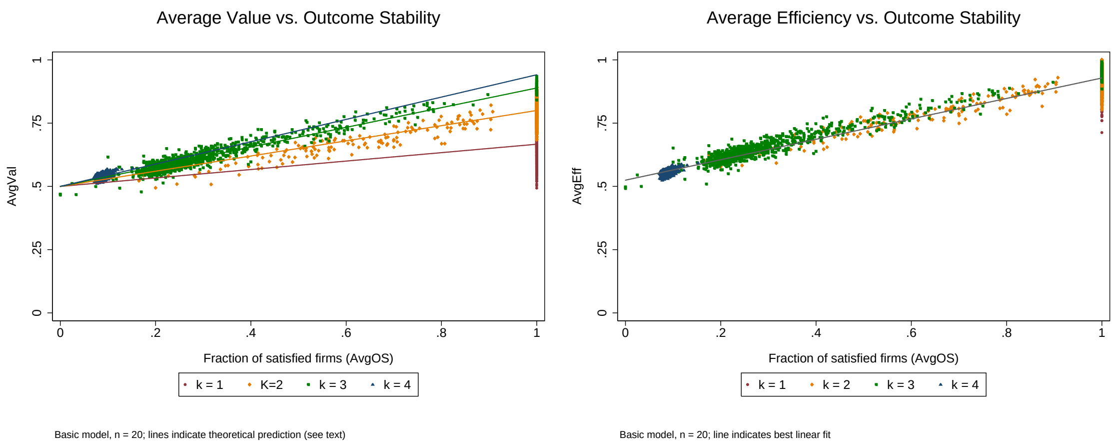

Figure 4: Average network value vs. outcome stability (left); efficiency vs. stability (right).

By our previous definition, outcome stability is simply the fraction of firms earning the best-response payoff versus the lower "frustration" payoff, so if our assumptions are correct, the average payoff for a given level of outcome stability should be a linear weighted average of 0.5 and $\frac{2^k}{2^k+1}$. The colored lines in the left-hand graph of Figure 4 plot these weighted averages for each value of k, which agree strongly with the data.

The right-hand graph of Figure 4 rescales the points on the left by dividing each value of AvgVal by the corresponding optimal value (as in Figure 3) to yield a plot of efficiency versus stability. It is startling to see that the points now appear to be scattered around a single line, plotted in gray, with correlation coefficient 0.97. This correlation supports the inference that in the basic model, stability and efficiency are closely aligned.

This is a subtle result because it arises from the interaction of two distinct causal mechanisms, one that determines the expected optimal network value (a statistical property of the random payoffs), and another that determines the relationship between outcome stability and average value (which itself arises through the interaction of the payoff structure and the best-response dynamics). Nonetheless, its interpretation is refreshingly straightforward: networks in which firms are able to achieve stable link configurations are likely to be more efficient than those in which firms are unable to realize their best-response payoffs and thus continually frustrated by the actions of their network neighbors.

3.3 The elusive inefficiency of lock-in

So far we have only considered steady-state averages of efficiency and stability. One might also inquire whether the firms' performance is affected by the amount of time they spend wandering around in the space of network states before settling into an equilibrium. In particular, perhaps firms suffer from locking into equilibria too early?

This conjecture is a natural one in light of the vast literature on exploration and exploitation (March 1991), an influential stream of which uses the NK framework to model organizational search (Levinthal 1997). In this work, rapid convergence is typically a sign of under-exploration or "lock-in" to inferior local optima. The best performance is typically found in trials where convergence occurs in later periods. Applying the same intuition to the model of this paper, one might predict that firms could achieve higher levels of efficiency by improving their collective ability to avoid premature lock-in while eventually settling into a stable state. Despite the appeal of this conjecture, there is little support for it in the experimental results.

For the 5,682 trials of Experiment 1 that converged to equilibrium, I computed an additional outcome variable, CycTick, indicating the period in which convergence occurred (i.e., $\text{OS}(t) = 1$ for all $t \geq$ CycTick). I plotted AvgEff against CycTick and ran a variety of linear regressions, including dummies for n and k and several functional transformations on CycTick. In most of these regressions the CycTick coefficient was insignificant at the 10% level and its magnitude was small; the sign was also inconsistent between regressions.

The explanation for this negative result lies in the fact that the firms are not engaged in collective search but rather co-evolutionary adaptation. As Kauffman (1993, p. 238) states:

> There is a fundamental difference between simple adaptive evolution and co-evolution. Evolution on a fixed fitness landscape ... is similar to the behavior of a physical system on a well-defined potential energy landscape. In both cases, the attractors of the "adaptive" process are local optima which are single points. In a coevolutionary process, however, the adaptive landscape of one actor heaves and deforms as the other actors make their own adaptive moves. Such coevolving systems may not in general have a potential function. Thus coevolving behavior is in no way limited to attaining point attractors which are local optima, nor is it even clear that coevolving systems must be optimizing anything whatsoever.

Thus, firms using only local decision rules cannot in general explore the space of network states effectively, nor can they necessarily exploit an efficient state when they find one. Yet, the second and third experiments show that firms with greater individual capabilities and richer institutional conditions were able to achieve substantially better performance—both individually and collectively—than firms in the basic model. We now turn to those results.

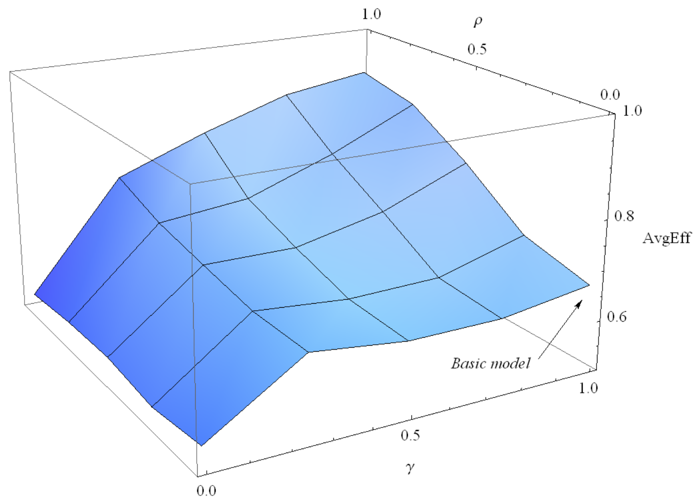

Figure 5: Economic efficiency as a function of perception (ρ) and cognition (γ).

4 Results: Industry and firm performance

This section presents the results of the two experiments that varied firms' individual capabilities and institutional conditions. As in the basic model, firms frequently became locked into inefficient outcomes or cycles of mutual frustration. However, simple coordinating institutions dramatically improved outcomes, even for firms with severely limited capabilities.

Like Experiment 1, Experiment 2 focuses on the industry level (i.e., the entire supply chain network), with economic efficiency (AvgEff) as the main outcome variable of interest. Experiment 3 introduces heterogeneity in firms' capabilities, and focuses on their performance at the individual level. Performance is measured in terms of cumulative payoffs, net of any transfers if the SidePayments institution is in effect.

4.1 Collective efficiency for firms without coordination

Figure 5 shows a graph of efficiency as a function of perception (ρ) and cognition (γ), under the Hobbesian institutional conditions with $n = 20$ and $k = 3$. As labeled, the front right-hand point ($\rho = 0$, $\gamma = 1$) corresponds to the basic model. Two features of the figure are immediately apparent. First, more of both parameters is generally better for efficiency. This finding affirms our basic intuition that more capable firms should be able to form more efficient networks. Second, the surface is quite smooth; for $\gamma > 0$ (i.e., firms that evaluate at least one action), small changes in the parameters yield correspondingly small changes in efficiency. This is consistent with the intuitive concepts of perception and cognition the model is intended to capture.

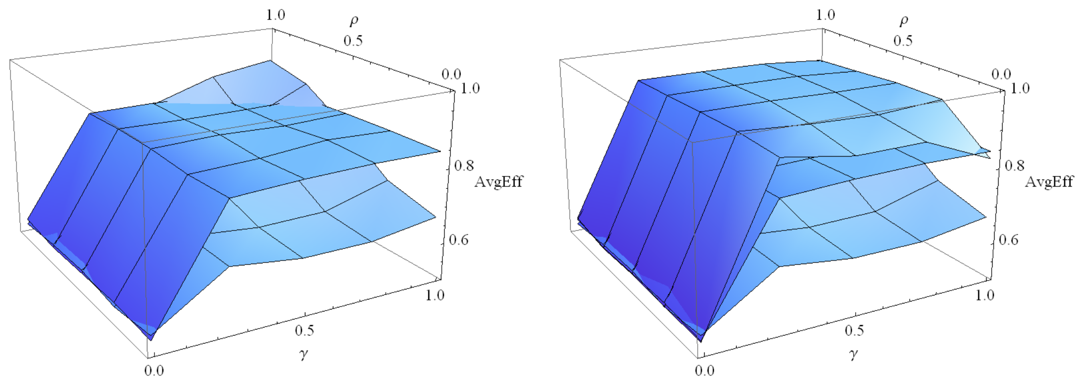

Figure 6: Economic efficiency, adding PropertyRights (left) and SidePayments (right).

Two less obvious features of the graph are also of interest. First, the slopes of the ridges converging to the peak of the surface ($\rho = 1$, $\gamma = 1$) are steeper than those near the origin. This indicates that perception and cognition are economic complements: more of one provides a greater boost to efficiency in the presence of the other. This observation is borne out in linear regressions, where the coefficient on an interaction term between the parameters is significant at the 0.1% level for all $k > 1$, with a magnitude that increases in k.

Second, the surface is kinked along the $\gamma = 0.25$ ridge. The kinked ridge also appears in the graph for $k = 4$, and presents a puzzle: why should low cognition—particularly in conjunction with low perception—yield higher efficiency than medium or high cognition? The intuition for the answer is that a firm with limited cognition can only evaluate a fraction of its potential actions in a given period, and if it fails to find an action with a higher expected payoff than the status quo, it chooses the latter. The result is a kind of inertia that benefits the firm's upstream neighbors: the lower the firm's cognition, the less likely it is to change its action in a given period, and in turn the lower its probability of frustrating its neighbors.

4.2 Collective efficiency with coordinating institutions

The left-hand graph of Figure 6 builds on Figure 5, overlaying a second surface that shows the average efficiency for trials in which the PropertyRights institution is in effect. The right side adds a third surface corresponding to the SidePayments case. Both graphs are for $n = 20$ and $k = 3$.

In the left-hand graph, we see another striking pattern: the top of the surface is almost completely flat. This is because the main effect of the PropertyRights institution is to drive the network dynamics to an equilibrium—not necessarily the efficient one, but at least a configuration in which each firm's link choices are a best response to the link choices of every other firm. In fact, all 2,500 trials represented by this surface converged to equilibrium, at a median period of $t = 4$. In the few periods before convergence, individual capabilities have little time to affect the efficiency of the network.

The other notable feature of the PropertyRights surface is the fact that its plateau occurs at an efficiency level of about 0.83, which is below the average efficiency achieved by firms with full perception and cognition under the Hobbesian "state of nature" (about 0.91). This can be seen by noting that the peak of the original surface pokes through at $\rho = 1$, $\gamma = 1$. The intuition for this result is that PropertyRights rules out a large set of network configurations, namely those with links that are perceived as harmful by the firms to which they are directed. The expected value of the best remaining individually rational configurations is lower than that of the best of the full set. The difference can be viewed as a cost of regulation. From a pure efficiency standpoint, the benefit that derives from increased stability outweighs the cost in the $k = 3$ case for all but the most capable firms. But for $k = 1$ and $k = 2$ (not shown), the PropertyRights rule strictly harms efficiency for all except zero-cognition firms. Property rights are thus beneficial for firms in high-complexity environments, but harmful when complexity is low.

Next, we turn to the right-hand graph of Figure 6. With the exception of the $\rho = 0$, $\gamma = 1$ case the ability to exchange side payments always yields higher average efficiency than simple property rights. This should not be surprising, since the SidePayments rule is conceptually just a stronger version of PropertyRights; links still must receive the consent of both parties, but no efficient states are ruled out. Interestingly, only about 60% of the trials converged to equilibrium under SidePayments. Even when firms cycled between several actions, though, they tended to avoid low-payoff configurations.

In the exceptional case of zero perception and full cognition, firms were "too smart for their own good." Able to evaluate their entire range of alternatives but basing their best-response calculations on outdated information, they often frustrated themselves by striking bargains that were not actually in their best interest. In this case, only 6% of the trials converged to equilibrium, and the firms would have been better off either under the PropertyRights condition or even the Hobbesian condition with full perception and cognition.

4.3 Individual profit for heterogeneous firms

While Experiment 2 focused on drivers of collective efficiency among firms with identical capabilities, Experiment 3 examined the individual profits of heterogeneous firms. The question in this context is whether perception and cognition confer similar benefits on firms in a

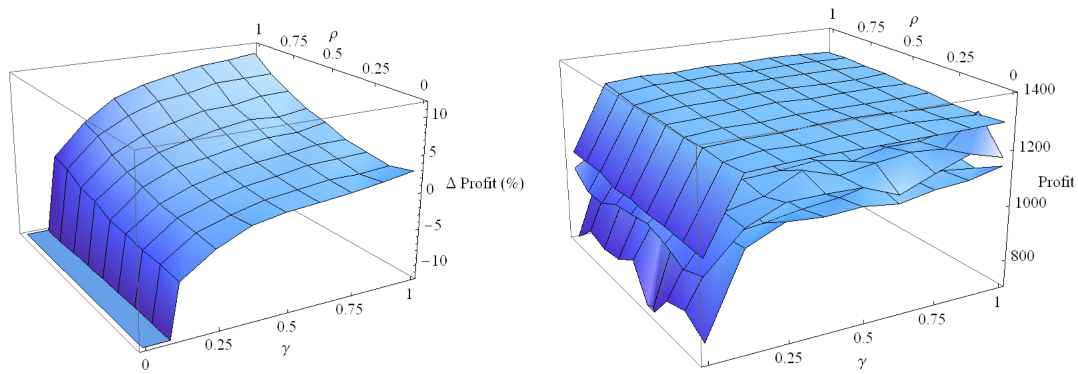

Figure 7: Difference in cumulative profit as a function of perception (ρ) and cognition (γ); cumulative profit under the Hobbesian, PropertyRights, and SidePayments conditions.

heterogeneous population compared to a population of identical firms. Figure 7 provides an affirmative answer.

The shape of the left-hand graph resembles that of Figure 5, so it is important to stress that what is being plotted is quite different. The vertical axis indicates the fraction by which, on average, the profit of a firm with perception and cognition given by the horizontal axes deviates from the average firm in its supply chain network. For example, firms with ρ and γ close to one earn on average about 10.6% more than the average among firms in the same experimental trial. The surface summarizes the differences in cumulative profit among the 200,000 firms in the 10,000 trials in which $k = 3$, $n = 20$, and the Hobbesian institutional conditions are in effect.

As in Figure 5, this surface generally slopes upward as the parameters increase, indicating that more capable firms earn higher profits when they are better endowed than their peers, just as the earlier figure showed that they form more efficient networks when all firms are identically endowed. And as in Experiment 2, regression results support the stronger claim that perception and cognition are economic complements with respect to firm profits. This finding indicates that firms have an incentive to invest in both kinds of capabilities; we could thus imagine a richer model in which these capabilities evolve endogenously.

As in Figure 6, the right side of Figure 7 plots a surface for each of the three institutional conditions. Here the values shown are cumulative firm profits instead of efficiency levels. Similar to the earlier results, the PropertyRights surface lies above the Hobbesian one, and the SidePayments surface lies above that. Again, both institutional innovations tend to neutralize the effects of capability differences among firms, yielding flatter surfaces than the Hobbesian case.

5 Discussion

The main finding of the paper may be summarized succinctly: local coordination mechanisms can be surprisingly effective, even when agents face complex incentive conflicts and even when their reasoning abilities are highly constrained. This section elaborates on the significance of this finding, the relevance of the model to real-world coordination problems, and the robustness and limitations of the experimental results.

5.1 Coase meets Simon, finds Hayek

The experiments showed that more capable firms tend to form more efficient networks, but even less capable firms can achieve high-value link configurations with the aid of simple local institutions that enable them to coordinate more effectively. It is not surprising that these institutions do not yield perfectly efficient outcomes. In the language of Coase (1960; 1988), firms in the model do not live in a world of zero transaction costs. In the language of Simon (1996), they are forced to "satisfice" by taking the actions that seem best given their limited perception and cognition. What is surprising is that these actions often lead to outcomes that are *close* to efficient, and that this happens not by solving a complex global coordination problem, but simply because the firms are able to achieve configurations in which neither they nor their immediate neighbors wish to change their behavior.

These results echo Hayek's observations about local coordination in the economy at large, which led him to marvel at the efficiency of the price system. In his celebrated paper on the use of knowledge in society, he evokes the image of a network in which "[t]he continuous flow of goods and services is maintained by constant deliberate adjustments, by new dispositions made every day in the light of circumstances not known the day before, by B stepping in at once when A fails to deliver" (Hayek 1945, p. 524). Although these networks may be complex, the fact that they tend to exhibit stable aggregate behavior cannot be understood in statistical terms alone. Instead, we need to look to the decentralized behavior of limited human agents, and ask what knowledge the "man on the spot" needs in order to make appropriate decisions about the resources within his control.

Hayek's principal answer, of course, was that prices serve as a signal of "the relative importance of the particular things with which he is concerned" (p. 525). Here there are no prices, however, and information about the relative value of particular link configurations does not literally propagate through the network. What *does* propagate through the network are the consequences of the firms' linking decisions, which can trigger cascades that undermine the stability of the entire supply chain. But the firms in the model benefit from the same kind of economic decoupling that Hayek identifies as a hallmark of a decentralized economy:

> There is hardly anything that happens anywhere in the world that *might* not have an effect on the decision he ought to make. But he need not know of these events as such, or of *all* their effects. It does not matter for him *why* at the particular moment more screws of one size than of another are wanted, *why* paper bags are more readily available than canvas bags, or *why* skilled labor, or particular machine tools, have for the moment become more difficult to acquire. ... The whole acts as one market, not because any of its members survey the whole field, but because their limited individual fields of vision sufficiently overlap so that through many intermediaries the relevant information is communicated to all. (pp. 525–526)

While the focus here is on the interaction between firms rather than the information they communicate, the same logic applies: because the consequences of their decisions propagate in a limited way, there is some hope that local coordination can be effective at inducing stability, and thereby yielding more efficient network configurations than those that arise through chaotic dynamics.

5.2 Scope and relevance of the model

In the model of this paper, firms are "smart enough" to seek locally optimal link configurations, but not so sophisticated as to render these configurations unstable by trying to influence firms other than their immediate neighbors to "jump" to a preferred peak in the network's fitness landscape. This observation sets important boundaries on the scope of the model with respect to real-world coordination problems. Specifically, the model does not provide useful insights about situations in which agents are either so "dumb" or "blind" as to be unable to even locally optimize their behavior, nor does it shed much light on situations in which agents engage in global optimization or multilateral negotiations. In these cases, existing models of (near-)random or (near-)rational behavior are more appropriate.

One might also ask whether the random payoff tables in the model are representative of real-world incentive structures. Many scholars have justified the application of Kauffman's NK framework to social and economic phenomena on the grounds that interactions between agents are often so complex that they are best modeled as random (Westhoff et al. 1996). The basic model of Section 2.1 is a member of the NKC family (Kauffman and Johnsen 1991), which was designed to study coevolution in biological systems but has since been applied to strategic behavior by firms (McKelvey 1999).[9] The use of random payoffs in the model should therefore not be controversial.

That said, the model is a "special case" in the sense that it assumes a strong form of payoff decoupling ($C = 1$, in the terminology of the NKC literature). Recall that a firm's payoff is independent of the choices of firms that are not its customers, and also of its customers' choices with respect to their other suppliers. It is easy to think of situations that violate this assumption, for example when a supplier of engine parts enters the market for transmissions, disrupting existing relationships between auto makers and their transmission suppliers. Nonetheless, the assumption captures a key feature of many loosely coupled systems (Glassman 1973; Weick 1976), namely the fact that—to paraphrase Hayek—even though everything may depend on everything else, we can often safely behave as if it doesn't. These kinds of situations are ubiquitous in social and economic settings (Orton and Weick 1990), and relate to what biologists call "weak linkage" (Kirschner and Gerhart 1998), which appears to be a fundamental property of evolvable systems.

5.3 Robustness and limitations

Even within the scope of local coordination by boundedly rational firms in loosely coupled systems, the model has many limitations. Probably the most significant is that while the firms interact repeatedly over time, their payoffs and decision rules are static. Extensions to the model could study both learning and environmental change using the same model setup. Many other decision rules and institutional conditions could also be explored, particularly with respect to bargaining (which is an active area of research in its own right). It would also be interesting to explore institutional heterogeneity, e.g., allowing some firms in a network to exchange side payments while others only to refuse links. These are all potential topics for future work.

In addition, while I conducted preliminary experiments to explore a variety of alternative assumptions (e.g., random and partially ordered topologies, starting at the optimal network configuration instead of the empty one, different interpretations of zero cognition and bargaining), any computational model can benefit from additional tests of its robust-

[9] Mapping the NKC parameters to the notation of this paper, the basic model is defined by $N = k$ (one "internal" element per supplier), $K = k-1$ (internal elements are fully interdependent), $C = 1$ (one "external" linkage per customer/supplier pair), $S = n$ (each firm is a "species"), and $X = k$ (each firm is linked to k suppliers). The only difference is that I do not distinguish the internal elements explicitly, which affects the shape of the payoff distributions. But most of the results in the NK framework do not depend sensitively on these distributions (Kauffman 1993).

ness to alternative parameter settings and implementation details. I did verify, though, that the main results are robust to the full set of parameters in the model ($n = \{10, 20, 30\}$, $k = \{1, 2, 3, 4\}$, and ρ and γ ranging from zero to one).

6 Conclusion

This paper explored strategic behavior in a network of firms using an agent-based model. The model exhibits a tension between economic efficiency and the stability of the network in the face of incentives to change its configuration. This tension is to be expected because the conditions of the Coase theorem are violated: the firms in the model lack the ability to discover efficient network configurations or achieve them through collective action.

Without property rights or the ability to exchange side payments, firms frequently became trapped in cycles of mutual frustration, unable to settle on a "peak" in their fitness landscape because these landscapes were repeatedly reshaped by their customers' choices (which were in turn reshaped by their customers, and so on). Although the institutional innovations studied in the extended model did not guarantee global efficiency—or even help the agents to "see the bigger picture" of their strategic situation—they led to better outcomes by favoring network configurations in which firms were locally satisfied. These results are consistent with Hayek's observations about the surprising effectiveness of local coordination in the economy at large, but the mechanism involved is different. Instead of prices, firms in the model exploited the fact that their incentives are partially decoupled by the structure of the network.

The results invite further study of the dynamic relationship between firms and their environment. In the current model, individual capabilities and institutional conditions are experimental parameters. In real life, however, we would expect them to coevolve. For example, in environments with weak institutions, we should find firms investing in perception and cognition to increase their performance. Similarly, in situations of high strategic complexity where even highly capable firms cannot perform well unaided, we can imagine institutions that facilitate coordination, like property rights and monetary exchange, to arise endogenously. Exploring these issues in a computational setting could help bridge the gap between existing models of network formation and the growing body of empirical work in this area.

References

Axtell, R. 2000. Effects of interaction topology and activation regime in several multi-agent systems. CSED Working Paper 12, Brookings Institution.

Binmore, K., A. Rubinstein, and A. Wolinsky. 1986. The Nash bargaining solution in economic modelling. *RAND Journal of Economics* 17(2):176–188.

Coase, R. H. 1960. The problem of social cost. *Journal of Law and Economics* 3(1):1–44.

———. 1988. Notes on the problem of social cost. In *The Firm, the Market, and the Law*. University of Chicago Press.

Derrida, B., and Y. Pomeau. 1986. Random networks of automata: A simple annealed approximation. *Europhysics Letters* 1(2):45–49.

Gershenson, C. 2004. Introduction to random boolean networks. In *Workshop and Tutorial Proceedings, Ninth International Conference on the Simulation and Synthesis of Living Systems (ALife IX)*, 160–173.

Glassman, R. B. 1973. Persistence and loose coupling in living systems. *Behavioral Science* 18(2): 83–98.

Goffman, E. 1974. *Frame Analysis: An Essay on the Organization of Experience*. Harper and Row.

Hayek, F. A. 1945. The use of knowledge in society. *American Economic Review* 35(4):519–530.

Iansiti, M., and R. Levien. 2004. *The Keystone Advantage: What the New Dynamics of Business Ecosystems Mean for Strategy, Innovation, and Sustainability*. Harvard Business School Press.

Jackson, M. O. 2004. A survey of models of network formation: Stability and efficiency. In *Group Formation in Economics: Networks, Clubs, and Coalitions*, ed. G. Demange and M. Wooders, 11–57. Cambridge University Press.

Jackson, M. O., and A. Wolinsky. 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71(1):44–74.

Kauffman, S. A. 1993. *The Origins of Order: Self-organization and Selection in Evolution*. Oxford University Press.

Kauffman, S. A., and S. Johnsen. 1991. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology* 149(4):467–505.

Kirschner, M., and J. Gerhart. 1998. Evolvability. *Proceedings of the National Academy of Sciences (USA)* 95:8420–8427.

Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science* 43(7):934–950.

Luke, S., C. Cioffi-Revilla, L. Panait, and K. Sullivan. 2004. MASON: A new multi-agent simulation toolkit. In *Proceedings of the Eighth Annual Swarm Users/Researchers Conference (SwarmFest 2004)*.

March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science* 2(1):71–87.

McKelvey, B. 1999. Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organization Science* 10(3):294–321.

Nash, J. 1953. Two-person cooperative games. *Econometrica* 21:128–140.

Nelson, R. R., and B. N. Sampat. 2001. Making sense of institutions as a factor shaping economic performance. *Journal of Economic Behavior and Organization* 44:31–54.

North, M. J., N. T. Collier, and J. R. Vos. 2006. Experiences creating three implementations of the Repast agent modeling toolkit. *ACM Transactions on Modeling and Computer Simulation* 16(1): 1–25.

Orton, J. D., and K. E. Weick. 1990. Loosely coupled systems: A reconceptualization. *Academy of Management Review* 15(2):203–223.

Osborne, M. J., and A. Rubinstein. 1994. *A Course in Game Theory*. MIT Press.

Press, K. 2007. When does defection pay? The stability of institutional arrangements in clusters. *Journal of Economic Interaction and Coordination* 2(1):67–84.

Rivkin, J. W. 2000. Imitation of complex strategies. *Management Science* 46(6):824–844.

Rivkin, J. W., and N. Siggelkow. 2002. Organizational sticking points on NK landscapes. *Complexity* 7(5):31–43.

Rubinstein, A. 1998. *Modeling Bounded Rationality*. MIT Press.

Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* 69(1): 99–118.

———. 1996. *The Sciences of the Artificial*. 3rd ed. MIT Press.

Stigler, G. J. 1966. *The Theory of Price*. 3rd ed. Macmillan.

Weick, K. E. 1976. Educational organizations as loosely coupled systems. *Administrative Science Quarterly* 21:1–19.

Westhoff, F. H., B. V. Yarbrough, and R. M. Yarbrough. 1996. Complexity, organization, and Stuart Kauffman's *The Origins of Order*. *Journal of Economic Behavior and Organization* 29(1):1–25.