2219

INTERNATIONAL POST DOCTORAL RESEARCH FELLOWSHIP PROGRAMME

FINAL REPORT

TITLE OF THE RESEARCH : Innovation complementarities in a network economy

RESEARCHER NAME: Cemil Faruk Durmaz

DIRECTOR NAME: Moritz Mueller

GENERAL INFORMATION

RESEARCH TITLE	Innovation complementarities in a network economy
RESEARCHER NAME	Cemil Faruk Durmaz
RESEARCH FIELD	Economics
RESEARCH DURATION	7 Months

Final Report Format:

1. Introduction

We model a network economy in which agents produce and trade goods and services. Traded goods include intermediate products, business services, or capital. Hence, the output of one agent may serve as input to production of another agent. A key point of the model is that each good comes with a certain quality (or technological level), and the quality of an input positively influences the quality of an output — how much depends on the capability of the producer. Because quality is valued by the consumer, producers have an incentive to produce high quality by buying high quality input as well as by investing in own quality related capability. The integration of intermediate goods over value chains thus induces innovation incentives along those chains. We are interested in endogenous system dynamics created by this mechanism. In particular under what conditions the system locks into a regime with low innovation efforts or high innovation efforts (with little and strong technological development respectively), and which policies and structural changes may lead to such regime shifts. In detail, we consider high and low-tech countries, the effect of integrated vs. separated networks, and closing vs. opening the economy to the world. Innovation policies include general vs. industrial specific rd subsidies, the effect of a state demand (in the military sector, space, or other), and the effect of varying trade tariffs. The effects of quality integration differ from other forms of innovational complementarities, or more general interdependence in innovation. More precisely, we differentiate this effect from innovation complementarity due to productivity improvements, which we model as cost reducing innovation but having no effect on the quality of the good (for theoretically distinguishing their effects rather than for them being separated in reality). We also allow for technology spillovers which are by definition non-traded (although potentially influenced by trade) and effective by providing knowledge useful for innovation (essentially the external pool of knowledge in Nathan and Rosenberg's chain linked model). Our focus on out of equilibrium dynamics leads us to run a simulation model. The model is not solved to derive an equilibrium solution1 , but used to investigate development paths of the system as a function of initial conditions, economic structure, exogenous events, and selection pressures. (We will need to take the most important of these.) For the beginning we consider myopic decision makers that make optimal decisions given the current system, but do not take into account future decisions by other agents (hence no common future vision can be developed by the agents - but this could be introduced later, I hope.) The number of agents in the model is fixed but unsuccessful agents may exit and successful agents may be replicated depending on the selection pressure of the economy.

One strain of literature we deal with centers on innovation complementarity: It is characteristic of a system that improvements in performance in one part are of limited significance without simultaneous improvements in other parts, just as the auditory benefits of a high-quality amplifier are lost when it is connected to a hi-fi set with a low-quality loudspeaker. (Rosenberg, 1979, p. 30) and on the same page: Similarly, improvements in power generation will have only a limited impact on the delivered cost of electricity until improvements are made in the transmission network and the cost of transporting electricity over long distances. (Rosenberg, 1979, p. 30) Thus, Rosenberg (1979) points out the existence of innovation complementarity stemming from integrating technology. The first example corresponds to what we term quality induced complementarity, the second is price induced complementarity. Bresnahan and Trajtenberg, 1995, in their GPT model, assume the existence of innovational complementarities, which they define as complementarity in the quality of the GPT and the quality of the application technology w.r.t. the private return in the application sector, in short $\partial \Pi a/(\partial gpt\ \partial a) > 0$. Another stream is on production or trade networks. Important papers are Hatfield et al., 2013 on trade networks, and Oberfield, 2018, on production networks. Both works center on stability and equilibria of exchange networks. In these works innovation is left as a footnote. Buera and Oberfield, 2020, model the diffusion of ideas - ignoring quality induced IC (innovation complementarities). Liu and Ma, 2023, propose a model of knowledge spillovers and investigate the central planner’s solution. Decision making of agents is not explicit in that model. Acemoglu, 2012, investigates how sectoral shocks lead to global shocks in a network economy (input-output matrices). This is also related - as it also applies to productivtiy shocks.

We present preliminary correlations in Figure 1 based on patent and research intensity weighted input-output tables (see data perspective section). Industries are grouped as Pavitt-taxonomy. We observe that supplier-dominated industries have positive correlation between externalities and domestic research, idea production. Science-based industries demonstrate positive correlation for patent production and patent related IO relations. Pavitt-taxonomy is essentially on categorizing industries based on their production differences. Thus, we expect to observe heterogeneity between agents i.e. firms in terms of supply chain and production technologies which is in line with our developed model as a motivating argument.

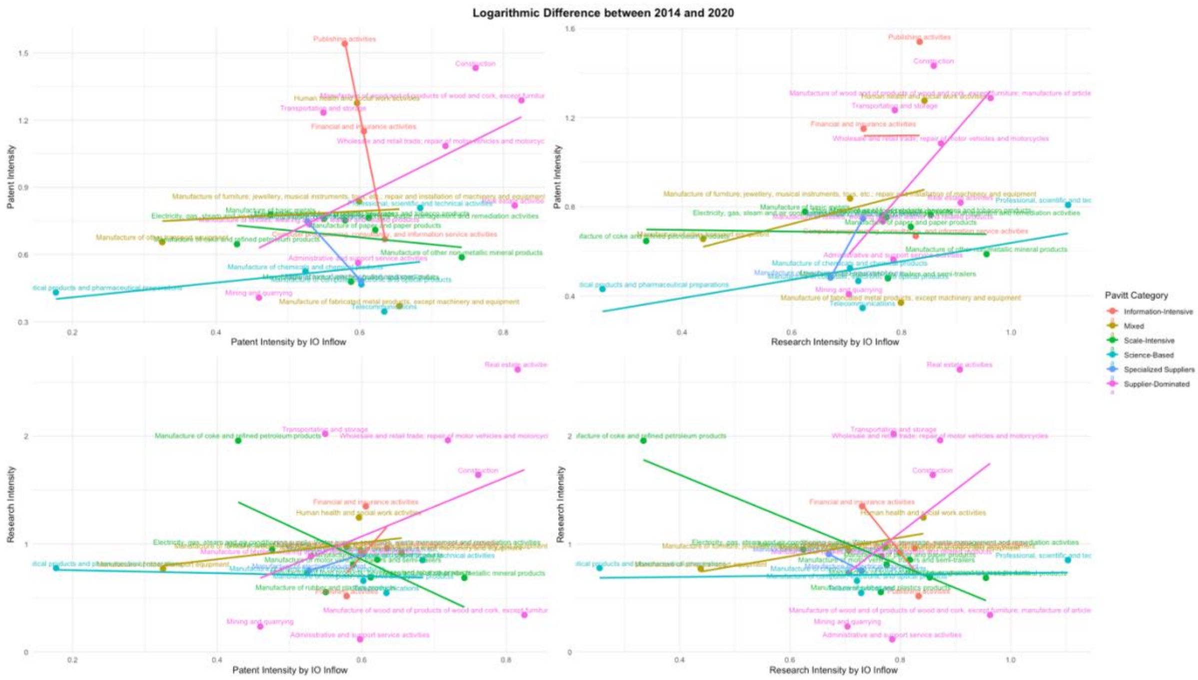

Figure 1. Research and Patents by IO Flows in Pavitt Taxonomy

2. Studies in report terms

I. Term: Developing the Model

A final consumer demands goods from various industries. Final demand is characterised by a constant elasticity of substitution (CES) utility function that aggregates quantity times quality over consumption goods. This has two implications. First, final goods produced in an industry are horizontally differentiated, that is they are only partly substitutable, which results in monopolistic competition in the industry. Second, consumers strictly trade-off quantity and quality such that final demand depends solely on the good's price per quality unit.

Households spend all income on consumption. Income stems from labor (inelastically supplied) for a fixed wage rate (our numeraire) and from profits.

Production takes place in a production network of interrelated tasks. A task is carried out by an agent in a certain industry. The agent transforms labor and (exactly) one intermediate good into one output, which can be sold to the final consumer or as intermediate good to another task in the same or another industry. Labor is a necessary factor in the production of quantities, but does not scale with quality. Both quality and marginal costs of production depend on capabilities of the producer in combination with the quality and marginal costs of its supplier. Agents invest in their capabilities through directed rd efforts, which are fixed sunk costs. And agents are free to choose their supplier.

Optimal pricing of final consumption goods with CES is achieved through a markup on marginal costs of production. Suppliers price intermediate goods by a variable cost component plus a fixed lump sum transfer. Variable cost must equal marginal costs of production to guarantee pairwise stability in supplier buyer relations. The additional lump sum transfer payment contributes to the agent's gross profit, which is used to cover rd sunk cost and, beyond that, contributes to agent's operating profit. Fixed transfer payments accrue from value chain profits: Any supplier producer transaction takes place within a value chain. Eventually, the most downstream agent sells with a markup to the consumer. This final transaction is the sole source of operational profit for the value chain (termed 'value chain profit'). Transfer payments within the value chain redistribute the value chain profit such that each agent gets a 'fair share', which we assume to be proportional to their relative contribution.

The economy is static. rd investment decisions and supplier choices of agents jointly determine the economy in terms of marginal costs of production, quality of goods, optimal pricing and profit allocation rules. In particular, the unique equilibrium where all markets clear and agents take optimal decisions is assumed to realise. Our main interest is in the rd efforts and supplier choices of profit oriented agents that anticipate the allocation in the resulting unique equilibrium.

Thus, the economy can be thought to realise in two steps. In a first step, agents decide on rd efforts and choose their suppliers. In a second step, the corresponding unique equilibrium realises. The model is analysed through simulation. The allocation in a realised unique equilibrium — conditional on rd efforts and supplier choices of the first step — is obtained numerically through fixed point iteration. The fixed point iteration thus maps agents' (rd and supplier) decisions on a unique profit allocation. This allows to construct a markov chain monte carlo (mcmc) sampler that generates a stable distribution of agents' decisions and resulting allocations. The key assumption in constructing the mcmc sampler is that agents are profit oriented: Given the decisions of all other agents, an agent is more likely to choose those actions that yield higher profits.

Set-up

Consider a static economy with $N = \{1, \dots, n\}$ agents partitioned over $M \geq 1$ industries, M_m. Each agent j produces exactly one good in a certain quality and quantity, (q_j, y_j). Production takes place in a production network. Necessary production factors are one intermediate good sourced from another agent i in a certain quality and quantity, (q_i, x_{ij}), and labor ℓ_i. Total output y_j is supplied to other agents for their production x_{jk}, or sold to a representative consumer in quantity d_j, hence $y_j = d_j + \sum_k x_{jk}$.

Demand

A representative consumer spends a total budget B in fixed proportions over markets to maximise total utility of consumption $U = \sum_m \beta_m \log(U_m)$, with $\sum_m \beta_m = 1$. Budget spent on a market is then $B_m = \beta_m B$. Within each market the consumer derives utility from consumption according to a CES utility function:

$$U_m = \left(\sum_{j \in \mathcal{M}_m} (q_j d_j)^{\frac{\sigma - 1}{\sigma}} \right)^{\frac{\sigma}{\sigma - 1}}$$

where $\sigma > 1$ is the elasticity of substitution. Given a price p_j and quality q_j for all goods in the market $j \in M_m$, the consumer maximises utility by demanding the quantity

$$d_j = \frac{q_j^{\sigma - 1} p_j^{-\sigma}}{\sum_{i \in \mathcal{M}_m} q_i^{\sigma - 1} p_i^{1 - \sigma}} \cdot B_m$$

Production technology

Production takes place in a network with agents N being the nodes, and supplier-buyer relations $E = \{ ij \mid i, j \in N \}$ the edges. Edges are directed as supplier i provides input to buyer j. There is at most one input to a task, whereas a task output may be delivered to multiple agents. The network of suppliers and buyers can therefore be described by a binary adjacency matrix A of dimension $(n \times n)$ with one entry $a_{ij} = 1$ in column j and all other entries zero. There is also a corresponding weighted adjacency matrix X with entries x_{ij} to capture the number of intermediate goods transferred among producing agents.

Two separate equations describe the quantity and quality of agent j's output:

$$y_j = \frac{1}{\alpha^\alpha (1-\alpha)^{1-\alpha}} \phi_j \, x_{ij}^\alpha \ell_j^{1-\alpha}$$

$$q_j = \delta \; \chi_j^\beta \; q_{i} + (1 - \delta) \; \psi_j$$

The quantity production function is Cobb-Douglas with productive efficiency φ_j, intermediate input x_{ij}, labor ℓ_j, and output elasticity of intermediate input $0 < \alpha < 1$.

Output quality, q_j, is a weighted average of two components; $0 \leq \delta \leq 1$. The first component is derived from the quality of the input good, q_i, in combination with the buyer's integrative capacity χ_j with exponent $\beta > 0$. The second component is autonomous quality ψ_j that is independent of the input.

A producing agent is thus characterised by three distinct forms of capabilities (φ, χ, ψ): productive efficiency φ, integrative capacity χ, and autonomous quality ψ. Capabilities are restricted to ensure that production in the network is feasible: minimal productive efficiency $\varphi_j > 1$; $0 < \chi, \psi < 1$ to bound quality $q \in (0, 1)$.

R&D investments

An agent may direct investments to each type of capability, $r = (r_\varphi, r_\chi, r_\psi)$. We refer to r as rd effort, but it may be interpreted as capturing any fixed, sunk costs directed at increasing productivity and quality (in the eye of the consumer). rd investments are mapped to firm capabilities in order to capture positive but decreasing marginal effects of rd, and to maintain the restrictions imposed on capabilities. In detail, we define production efficiency as $\varphi = r_\varphi^\kappa + 1$. On the quality side, integrative capacity is calculated as $\chi = r_\chi^\kappa/(1 + r_\chi^\kappa)$; the same for autonomous quality ψ. To keep the model simple, $\kappa < 1$ is a parameter common to all capabilities.

Prices and profits

Final demand induces economic activity in the production network along a set of value chains, Ω. And any economic activity can be attributed to a specific value chain, say $\omega \in \Omega$. The transaction with the consumer marks the end of a chain ω, and is the source of value chain profits. Let m_j be agent j's marginal cost of production. If agent j sells to a consumer, he prices the consumer by a markup on marginal costs, $p_j = \sigma/(\sigma-1) \cdot m_j$. This price maximises the operating profits of the transaction, realising a value chain profit of $\pi(\omega) = 1/(\sigma - 1) \cdot m_j d_j$.

In chain ω, supplier i provides buyer j with a quantity $x_{ij}(\omega)$ that complies with final demand of that chain. The same supplier buyer relationship may serve multiple value chains, $\sum_\omega x_{ij}(\omega) = x_{ij}$, but assume a separate contract exists for each value chain. Such a contract would specify units sold in quantity and quality ($x_{ij}(\omega)$, q_i) and a payment that specifies a unit price and a fixed lump sum payment for the overall transaction, say ($p_i(\omega)$, $f_{i}(\omega)$).

The unit price equals the supplier's marginal costs, i.e. $p_i(\omega) = m_i$. This can be derived in two steps following (Oberfield, 2018). The first step is to derive buyer j's marginal costs as a function of supplier price. This yields the first equation:

$$\begin{aligned} m_j &= \frac{\partial c_j}{\partial x_j} \\ &= \lambda_j^{-1} w^{1-\alpha} p_{i}(\omega)^\alpha \\ &= \lambda_j^{-1} w^{1-\alpha} m_i^\alpha \end{aligned}$$

The second equation then follows from the observation that bilateral gains to trade are maximised only when the shadow value that buyer j places on input i, i.e. $p_{i}(\omega)$, is equal to supplier i's marginal costs m_i.

The fixed price component $f_i(\omega)$ is, in principle, subject to bargaining between the contracting partners. As marginal costs are payed by the variable price component, any bargaining concerns the allocation of value chain profit $\pi(\omega)$. In the model, we allocate value chain profits proportional to the agents' relative contribution i.t.o. rd efforts, assuming the value chain is fixed:

$$f_j(\omega) = \frac{\Delta \pi_j(\omega)}{\sum_i \Delta \pi_i(\omega)} \pi(\omega)$$

Thus, the profit of agent j in the production network becomes:

$$\pi_j = \sum_{\omega \in \Omega_j} f_j(\omega) - \sum_{x \in \{\phi,\chi,\psi\}} r_{j,x}$$

All economic activity in the production network is induced by final demand.

The final transaction with the consumer marks the end of a value chain, say ω.

Subsequent agents in the value chain ω have supplier buyer relationships. A supplier i provides buyer j with a quantity $x_{ij}(\omega)$ that complies with final demand of that chain.

Note that all activity in the production network is associated with a particular value chain that is induced by a final transaction, and the same supplier buyer relationship may serve multiple value chains, i.e. $\sum_\omega x_{ij}(\omega) = x_{ij}$. Assume that agents

transact intermediate goods in a quantity $x_{ij}(\omega)$ that matches final demand of that value chain.

Note that all activity in the production network is associated with a particular value chain, and the same supplier buyer relationship may serve multiple value chains, i.e. $\sum_\omega x_{ij}(\omega) = x_{ij}$.

transact intermediate goods from

Later we see that agent j only keeps a fraction of this profit but that does not change the optimisation result.

The agent maximises operating profits from a transaction with

An agent selling to the final consumer is by construction the most downstream firm of a value chain, say ω. In the value chain, upstream firms

contribute through their supply to the final product.

Any production takes place in a value chain, and all value chains end with the most downstream firm selling to the consumer. Thus all economic activity in the production network can be attributed to a

If a product is sold to the consumer a value chain ω ends.

Firms set consumer prices as a markup on marginal costs, i.e. $p_j = \sigma/(\sigma - 1)\, m_j$. This is the profit maximising price given CES demand.

Contracts among firms specify units sold (x_{ij}, q_i) and the payment which includes a variable cost and a fixed cost component. Variable costs are set equal to marginal costs to ensure pairwise stability in a setting that allows for transfer payments. If transfer payments are allowed, rational contract partners maximise their joint value of the transaction and the variable price equals marginal cost rule ensures alignment to that common goal. More precisely, the price set by the supplier must be equal to the supplier's marginal cost and match the shadow value the buyer places on the input (see Oberfield, 2018). This allows to formulate the marginal costs of the buyer m_j as a function of the supplier's marginal price m_i:

$$m_j = \frac{\partial c_j}{\partial x_j} = \lambda_j^{-1} w^{1-\alpha} m_i^\alpha$$

A fixed cost component is necessary to recover sunk costs (here rd). How much is, in principle, subject to bargaining between the contracting partners. The profit that is bargained over arises at the end of a value chain, when the final product is sold to the consumer with a markup over marginal costs; call that the value chain profits. All agents within the value chain somehow contribute to the value chain profit and therefore may demand their 'fair' share. We consider rd sunk costs of value chain members as a form of effort that contributes to the value chain profits. The marginal contribution of agent's j rd efforts to value chain profits is calculated as the difference between the actual value chain profits and the value chain profits in the same configuration but with agent j doing no rd at all. Profits are then distributed across members of the value chain proportional to their marginal contribution.

Profits The profit of agent j is calculated as follows:

$$\pi_j(\omega) =$$

$$\pi_j = p_j d_j + \sum_{\omega} (p_j x_{jb(\omega)} + f_\omega) - p_i x_{ij} - w l_j - \sum_x r_{x,j}$$

Clearly, profits depend on the actions of the focal agent, a_j, and actions of all other agents a_{-j}, which we emphasize by writing $\pi_j(a_j, a_{-j})$.

Thus, an agent is endowed with a vector of capabilities $\lambda = (\lambda_x, \lambda_{(q,ij)}, \lambda_{(q,i)})$. The different capabilities are effective in different ways: λ_x defines production efficiency, $\lambda_{(q,ij)}$ is the capability to complement or substitute the quality of a specific input, and $\lambda_{(q,j)}$ defines the quality part of the produce that is independent from the input.

Each agent may direct investments in each type of capability, $r = \{r_{x}, r_{(q,ij)}, r_{(q,i)}\}$, which we refer to as R&D effort but in reality may also capture any other effort directed at increasing productivity or quality (in the eye of the consumer).

We assume $\lambda_x = r_x^\kappa$ with $\kappa < 1$ to capture decreasing marginal effects. Capacities related to qualities are further constrained to range between zero and one: $\lambda = r_x^\kappa/(1 + r_x^\kappa)$ for $\lambda_{(q,ij)}$ and $\lambda_{(q,ij)}$.

TODO complementarity and subsitutability of integration efforts and input quality. ... But maybe not so important because in the end complementarity is w.r.t. payoffs.

First note that the production network may exhibit cycles. Consider production of quantities. If production efficiencies φ within a cycle are too small, agents need more material input than they produce as output. In a cycle where the output of each agent serves as input to the next agent in the cycle, φ too small will lead to infinite demand of intermediate input. Therefore we restrict $\varphi > 1$, that is we impose a minimal production efficiency.

In quality production the opposite is the case. Intuitively, if each agent in a cycle improves over the quality of the supplier quality becomes infinite. Applying the restriction on integrative capacity $0 < \chi < 1$ and on autonomous quality $0 < \psi < 1$ ensures that qualities range from zero to one. We consider this as a qualitative scale. A quality of 0 is the most basic product in the market, and a quality of 1 corresponds to the maximum quality that can be possibly attained. Note that in the demand function, zero qualities are avoided by rescaling $(q + 1) \cdot 10$ with a resulting range of $q \in [1, 11]$.

Feasibility of production in quantities

We first note that each agent produces total demand y_j at minimal costs. Given wage rate w and input price p_i, the cost-minimizing input ratio in the Cobb-Douglas production function is $x_{ij}/l_j = (\alpha/(1 - \alpha)) (w/p_i)$.

This implies that total production can be expressed as a function of only supply without labor:

$$x_j = \frac{1}{\alpha} \lambda_{xj} \left(\frac{p_i}{w} \right)^{1-\alpha} x_{ij}$$

In addition, the interaction structure A ($A_{ij} = 1$ if ij) can be used to express the production system as follows:

$$\boldsymbol{x} = \boldsymbol{X} \mathbf{1} + \boldsymbol{c}$$

$$\boldsymbol{x} = \boldsymbol{M} \boldsymbol{x} + \boldsymbol{c}$$

$$\boldsymbol{x} = \left(\mathbf{I} - \boldsymbol{M} \right)^{-1} \boldsymbol{c}$$

where X is the input-output matrix with entries x_{ij}, $M = \mathrm{diag}(\omega)^{-1} A \,\mathrm{diag}(\lambda_x)^{-1}$, with $\omega_i = (1/\alpha) (p_i/(w^{*}))^{1 - \alpha}$.

The matrix $(I - M)$ is invertible if the spectral radius $\rho(M)$ is less then 1:

$$\rho(M) = \max \left\{ |\lambda| : \lambda \in \operatorname{eig}(M) \right\} < 1$$

For the interaction matrix $\rho(A) \leq 1$ with one if the interaction matrix includes a cycle.

The intuition: efficiencies or incentives to augment labor (relative to supplies) must be high enough to prevent the need of infinite production of supplies (in a cycle).

Feasibility of production

Both components of quality, and hence also output quality, are designed to range between 0 and 1. The value 0 can be interpreted as a quality that is too low to draw any utility from the output (say moldy bread). The value 1 is a level of quality that corresponds to the maximum potential quality one can possibly imagine (the softest bread in England, or the crunchiest in Germany).

Without such a restriction, quality 'production' becomes infeasible through circular quality improvements.

Agent j offers that product for a price

p_j and for that price produces the amount x_j as demanded by the market. Two separate equations describe the mapping of quantities (x_j) and qualities (q_j):

$$x_j = \frac{1}{\alpha^\alpha (1 - \alpha)^{1 - \alpha}} \lambda_{j}^{\text{prod}} x_{ij}^\alpha l(q_j)^{1 - \alpha}$$

$$q_j = \delta \; \lambda_{ij}^\beta \; q_{i} + (1 - \delta) \; \lambda_j$$

The quantity production function is standard — a Cobb-Douglas function with input x_{ij} and labor $l(q_j)$.

Agents in each industry use intermediate goods and labor to produce intermediate goods or final goods, or both. Agents are endowed with a production technology and are restricted purchase their supply from a certain industry. The decision of each agent to invest in rd, sunk costs that determine its capability in the production of quantities and qualities, and their choice of a specific supplier generates a unique allocation of the economy.

Consumers provide inelastically labor L for the production of goods for which they receive a common wage rate w, yielding a budget $w L = C$ to spend on goods and demand: [Footnote: Alternatively profits P of producing agents could also be consumed, yielding total consumption expenditure $C = I + P$. Subsequent sections may elaborate the demand side. We first consider a closed economy in one country. The economy is partitioned into a number of sectors or industries M_m, $m = 0, \ldots, M$. Sector M_M is a consumption sector that exerts final demand and hence produces nothing. A representative consumer values quality and quantity from different industries. We define the consumer's utility as a two-level nested CES aggregator. Let C_m be the aggregate consumption index for industry m. The total utility is $U = (\sum_{m=1}^{M} C_m^{(\sigma_o - 1)/\sigma_o})^{\sigma_o/(\sigma_o - 1)}$ where $\sigma_o > 1$ is the elasticity of substitution across industries. Each C_m is itself a CES aggregator of firm-level consumption x_i, adjusted for quality q_i: $C_m = (\sum_{i \in m} (q_i x_i)^{(\sigma_i - 1)/\sigma_i})^{\sigma_i/(\sigma_i - 1)}$ where $\sigma_i > 1$ is the elasticity of substitution within industry m.]

Investments in capability

Thus, an agent is endowed with a vector of capabilities $\lambda = (\lambda_x, \lambda_{(q,ij)}, \lambda_{(q,i)})$. The different capabilities are effective in different ways: λ_x defines production efficiency, $\lambda_{(q,ij)}$ is the capability to complement or substitute the quality of a specific input, and $\lambda_{(q,j)}$ defines the quality part of the produce that is independent from the input.

Each agent may direct investments in each type of capability, $r = \{r_{x}, r_{(q,ij)}, r_{(q,i)}\}$, which we refer to as R&D effort but in reality may also capture any other effort directed at increasing productivity or quality (in the eye of the consumer).

We assume $\lambda_x = r_x^\kappa$ with $\kappa < 1$ to capture the idea that increasing efforts always increases capability but do so with decreasing marginal effects. Capacities related to qualities are further constrained to range between zero and one by assuming $\lambda = r_x^\kappa/(1 + r_x^\kappa)$ for $\lambda_{(q,ij)}$ and $\lambda_{(q,ij)}$.

TODO complementarity and subsitutability of integration efforts and input quality. ... But maybe not so important because in the end complementarity is w.r.t. payoffs.

Agent behaviour and (aggregate) outcomes

Each agent j is characterized by $(r_j, \lambda_j, s_j, (q_i, p_i, x_i), l_j, (q_j, p_j, x_j))$, i.e. rd expenditures, capabilities, supplier, input characteristics, labor hired, and output characteristics respectively. The basic model offers three decision variables. Choice of supplier s_j, rd expenditures r_j, and price offer p_j together define the agent's action $a_j = (s_j, r_j, p_j)$.

Before proceeding to how agents make their choices, we describe how the actions of all agents in combination $a = (a_1, \dots, a_n)$ define one realisation of the economy X out of all possible realisations $\mathcal{X}$. The space of possible realisations of the economy is restricted by boundary conditions and feasibility (or consistency).

Supplier choice, $s_j = (i; i \in S_j)$, is a discrete choice among (exogenously fixed) potential suppliers S_j. Supplier choices of agents $s = (s_1, \dots, s_n)$ jointly generate a directed binary network $G \in \mathcal{G}$, where $\mathcal{G}$ is the set of all potential networks.

The basic model is static, i.e. without time. Therefore, rd expenditures immediately result in capabilities for all agents, i.e. $\lambda_j = r_j^\kappa$. Within a supplier network G, agents apply their capabilities Λ to achieve a certain product quality $(q_1, \dots, q_n)$. This can be calculated numerically under conditions that guarantee a contractive mapping as shown below in subsection ‘consistent quality production’.

When prices $p = (p_1, \dots, p_n)$ are set, we assume agents carry out all market orders. Market orders x_j consist of quantities ordered from final demand c_j and intermediate demand x_{jk}, i.e. $x_j = c_j + x_{jk}$. Production is consistent (see ‘consistent quantity production’), and implies labor hired $\ell = (\ell_1, \dots, \ell_n)$.

The profit of agent j is calculated as follows:

$$\pi_j = p_j c_j + \sum_{k} p_j x_{jk} - p_i x_{ij} - w l_j - \sum_x r_{x,j}$$

Clearly, profits depend on the actions of the focal agent, a_j, and actions of all other agents a_{-j}, which we emphasize by writing $\pi_j(a_j, a_{-j})$.

Feasibility of qualities

Only rd decisions r_j affect quality q_j, pricing decisions do not. Output quality of one agent depends on output quality of other agents. Because that relationship is non-linear, we find iteratively a solution through contractive mapping. Let $q^{(t)}$ be the current guess at iteration t. Then the update rule is:

$$q_j^{(t+1)} = \delta \lambda_{q,ij}^\beta q_i^{(t)} + (1 - \delta) \lambda_{q,j}$$

This is a bounded, smooth function of q_i. The derivative $\partial q_j / \partial q_i = \delta \lambda_{q,ij}^{\beta}$ with δ, $\lambda_{q,ij} < 1$ shows it is a contraction. Furthermore $\partial^2 q_j /(\partial q_i \partial \lambda_{q,ij}) = \delta \beta \lambda_{q,ij}^{\beta-1} > 0$ ($\delta, \beta, \lambda_{q,ij} > 0$), i.e., input quality and integrating efforts are complementary w.r.t. output quality.

Feasibility of quantities

Production is consistent if i) all available labor is used for a wage rate $w*$, ii) consumers maximise utility given their budget implied by that wage rate (or wage rate and profits), iii) total production equals intermediate plus final demand for all producing agents, iv) no coalition of agents has an incentive to deviate from the production schedule by proposing a different price or changing supply.

Cost-minimising input ratio in Cobb–Douglas is $x_{ij}/l_j = (\alpha/(1 - \alpha)) (w/p_i)$. This implies:

$$x_j = \frac{1}{\alpha} \lambda_{xj} \left(\frac{p_i}{w} \right)^{1 - \alpha} x_{ij}$$

System representation:

$$\boldsymbol{x} = \boldsymbol{X} \mathbf{1} + \boldsymbol{c}$$

$$\boldsymbol{x} = \boldsymbol{M} \boldsymbol{x} + \boldsymbol{c}$$

$$\boldsymbol{x} = \left(\mathbf{I} - \boldsymbol{M} \right)^{-1} \boldsymbol{c}$$

with $M = \mathrm{diag}(\omega)^{-1} A\ \mathrm{diag}(\lambda_x)^{-1}$, $\omega_i = (1/\alpha) (p_i/(w^{*}))^{1 - \alpha}$. Invertibility condition (finite requirements):

$$\rho(M) = \max \{ |\lambda| : \lambda \in \operatorname{eig}(M) \} < 1$$

For the interaction matrix $\rho(A) \leq 1$ with one if the interaction matrix includes a cycle. Efficiencies/incentives must be high enough to prevent infinite intermediate needs in cycles.

Short discussion of the basic model set-up

The theory is exposed in a certain language in order to remain clear and to relate to discussions in the current literature. There is however a clear relation to more traditional economic terminology. What we call quality could be also termed product embedded technology. Capability could be also knowledge or technology possessed by the agent, etc.

The two components of quality enter as a weighted sum with parameter δ ($0 \leq \delta \leq 1$) to reflect the fact that output quality may depend more or less on input quality. If the task is repackaging in wholesales for example, input quality will largely determine output quality (δ close to one). In contrast, if the task is to make a drawing with a pencil, the pencil's quality only modestly influences the artwork's quality; corresponding to δ close to zero.

Discussion of unlimited quality due to cycles in the network where the geometric average of quality feedbacks exceeds one.

This takes up ideas from the literature on technologies and recipes, and on production networks. Formally it can be generalized to multiple sectors as follows: [Footnote elaborating multi-sector generalization …]

Consider a stone age economy involving two agents, Robinson and Friday: … Bounding quality by a potential maximum of one ensures defined quality even in cycles in the task network.

Consistent with interpreting the production network as a task network, one may impose acyclicity. But that would limit the model to product integration chains, only a part of value chains.

Profits are then distributed across members of the value chain proportional to their marginal contribution.

II. Term: Data Perspective

We collect data for the United States from the OECD database and deflate all observations with the sector gross output price index (GOPI) from the KLEMS database as follows, where X is the value in a certain industry for a year and P is the price index with 2015 as the base year.

$$X^{real}_{it} = \frac{X_{it}}{P_{it}}100$$

R&D

We collect data for Business enterprise R&D expenditure by industry (BERD) from the OECD database in real dollars and calculate R&D capital stock using perpetual inventory with an arbitrary depreciation rate δ for each industry and year as follows:

$$r_{it} = R_{it} - \delta R_{i(t-1)}$$

Mark-up (Profit-Cost Ratio)

We calculate the mark-up for each industry based on Castelo and Puty (2018) as follows:

$$MU_{it} = \frac{\Pi_{it}}{W_{it} + I_{it}}$$

Gross Operating Surplus is the proxy for industry profit, and W stands for labour compensation, I is intermediate good purchases from other industries in the Input-Output tables. The calculation of mark-up does not present price over costs, but rather the profit-to-cost ratio. So it is possible that the ratio is less than 1, referring to a lower share of profits in the total revenue of the industry compared to costs, but still positive profits.

Ideas and Knowledge Stock

We measure ideas produced in each industry by using number of patents from OECD. The database provides number of patent applications in 35 WIPO technology fields to WIPO as Patent Authority with application date of the inventor. We use probabilistic concordance schemes (Neuhäusler, P., Frietsch, R., and Kroll, H., 2019) to disperse technology fields into Nace rev. 2 sectors. We also calculate knowledge stock with the perpetual inventory method as follows:

$$k_{it} = K_{it} - \delta K_{i(t-1)}$$

Inter Industry Transactions

The literature considers input-output tables as representation of an economy's industrial structure where each industry is connected to others through transactions of goods and services bilaterally. We utilize national IO tables from OECD to observe transactions between industries in real terms and calculate purchases of each industry from others as well as final demand to each industry as end-use products.

In Figure 1.1, edges stand for intermediate good transactions and nodes are patent application by each industry. The plot suggests that industries such as pharmaceuticals, which are highly patenting, do not necessarily connect to other industries in IO as much as construction.

We also observe that industry mark-up and research expenditures do not necessarily overlap with the transaction intensity between industries in the IO table in Figures 1.2 and 1.3. Thus, we expect that for the relationship between ideas, profitability, research efforts and inter-industry sales, heterogeneity may arise from all which requires a multi-layer network approach to disintegrate the structure.

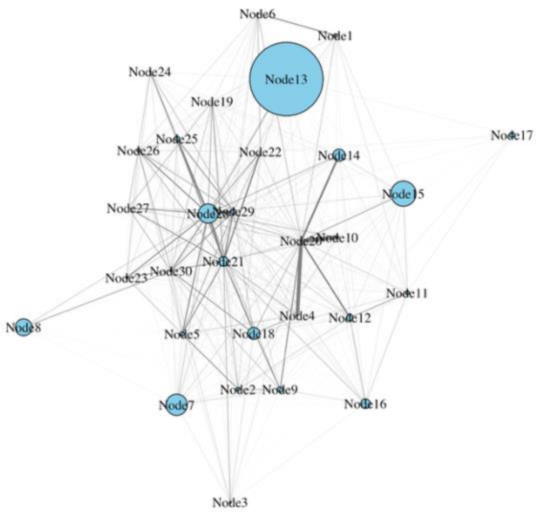

[Figure 1.1 IO transactions and Patent applications]

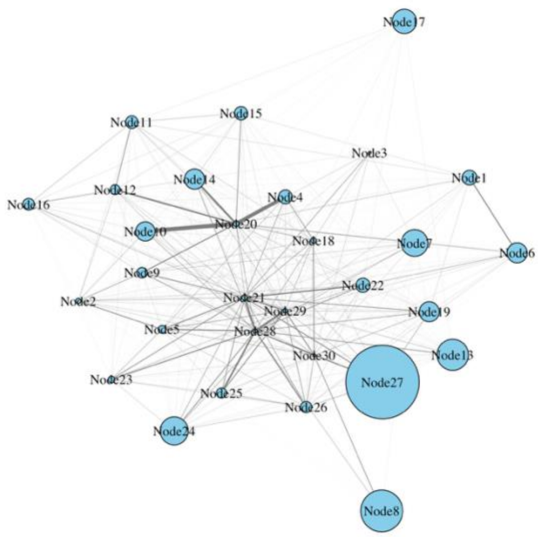

[Figure 1.2 IO transactions and Mark-up]

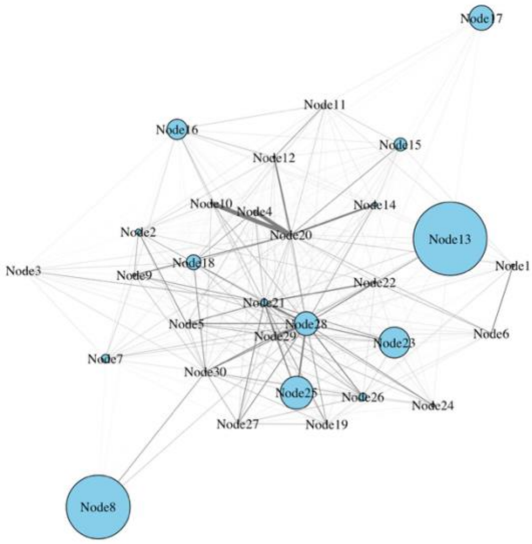

[Figure 1.3 IO transactions and Research]

However, in Figure 2, we observe that markups are positively correlated with patents and R&D, and negatively correlated with IO transactions (purchases and sales of intermediate goods, final demand of consumers). Patents and research are in a positive relationship with IO transactions. We expect industries to trade with others who do more research and have high research and knowledge, correlated with high profitability. However, high markup does not seem to be related to transactions between industries; rather, it is negatively correlated.

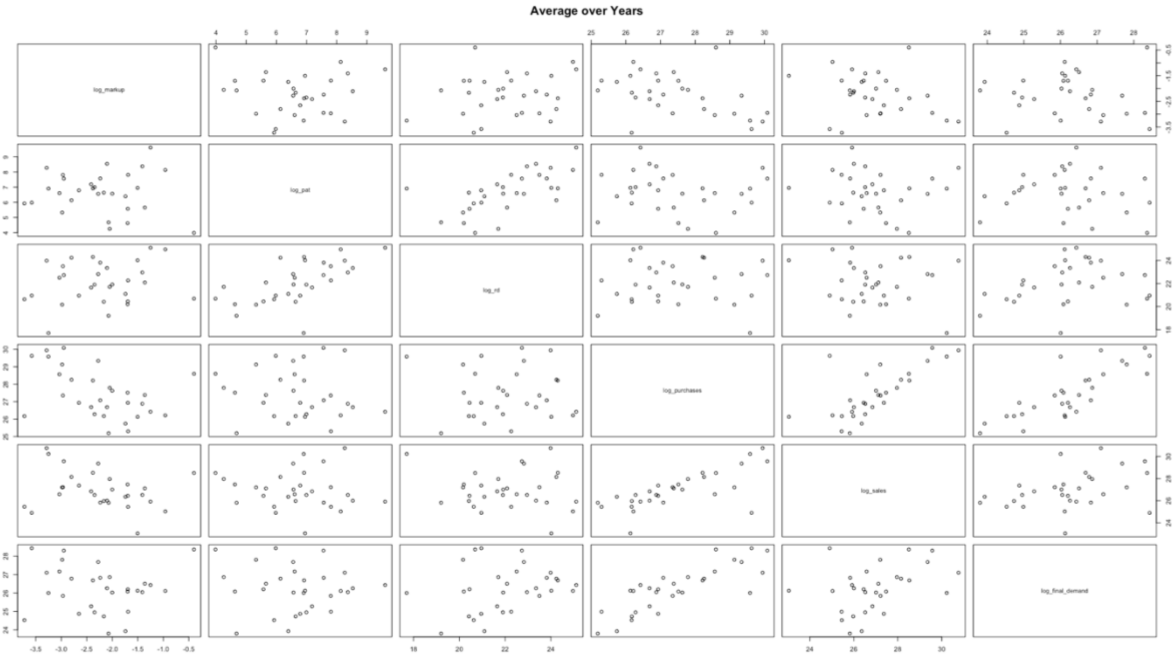

[Figure 2: Correlations over the years]

We interact IO matrix with research expenditure and patents separately in Figure 3 and observe highly positive relationships between patent applications by industries and research, patents they receive from other industries through input purchases.

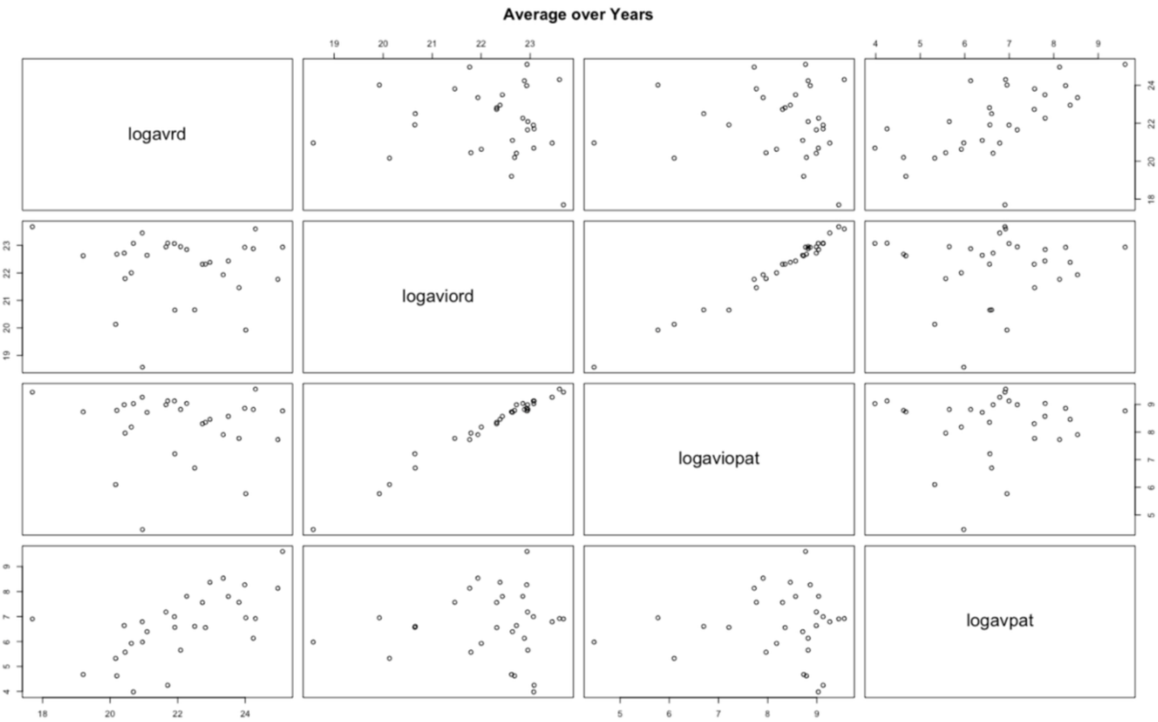

[Figure 3: IORD Correlations over the years]

Technology Spillovers

We calculate technological structure between industries based on Probabilistic concordance schemes (Neuhäusler P., Frietsch R. and Kroll H., 2019) as we did for the patent data as follows where w stands for technology field and i is the industry. We multiply the concordance table, i.e., adjacency matrix A_{wi}, by the number of patent applications in each tech field and row-normalise, then apply cosine similarity to produce patent shares between industries.

$$q_{wi} = q_w A_{wi}$$

$$\tilde{q}_{wi} = \frac{q_{wi}}{\sum_{j=1}^{n} q_{wi}}$$

$$q^{norm}_{wi} = \frac{\tilde q_{wi}}{\sqrt{\sum_{j=1}^{N} \tilde q_{wi}^2}}$$

$${q}_{ji} = q^{norm}_{wi}(q^{norm}_{wi})^T$$

Finally, we multiply the sum of each column by R&D expenditure of industry i to calculate R_i.

$${q}_{ji} = \tilde q_{wi}(\tilde q_{wi})^T$$

Thus, the total spillover received by industry i interacted with the research efforts of other industries is denoted by:

$$T_i = r_j \sum_{j=1}^{N} q_{ji}$$

We observe that residuals from the correlation between patents and R&D expenditures present correlation heterogeneity for different industries in the US under the Pavitt taxonomy classification. Thus, Figure 4 shows that an unexplained part of the variation in patent applications is the result of technology spillovers. It can be seen that this relationship seems positive for the supplier-dominated and information-intensive sectors.

We consider that the technology spillovers are potentially from science and information-intensive industries to other industries, stronger and directed. We also observe that pharmaceuticals benefit least from spillovers from other industries, and it is plausible that due to the nature of the industry, pharmaceuticals are a closed, intensively patented industry which produces knowledge with in-house R&D.

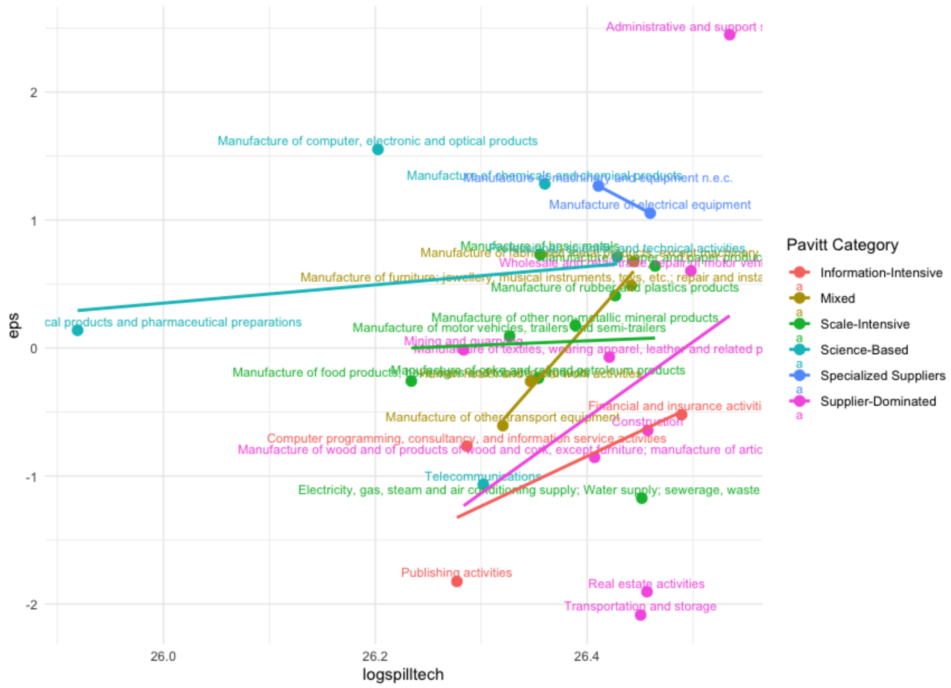

[Figure 4: log(T_i), ε against patent applications]

III. Term: Analysis of the Model

Network and R&D Decisions Given

We show in the following that a unique equilibrium is attained when network and R&D decisions are given.

R&D and Network

The analysis investigates likely realizations of our economy under the assumption that agents prefer actions that yield higher profits. Under this assumption, we specify the conditional distribution of a_j as follows (we adopt the notation for discrete choice in the following but the reasoning applies to any choice):

$$
\begin{aligned}
p(a_j \mid a_{-j},\bm{\theta}) &= \frac{\exp(\alpha \pi_j(a_j, a_{-j}; \bm{\theta}))}{\sum_{a'_j \in \mathcal{A}_j} \exp(\alpha \pi_j(a'_j, a_{-j}; \bm{\theta}))} \\
&= \frac{\exp(\alpha \pi_j(a_j, a_{-j}; \bm{\theta}))}{c_j(a_{-j}; \theta)}
\end{aligned}
$$

The second equation notes that the denominator of this conditional probability is a normalizing constant in that it integrates out all potential actions a'_j.

Discussion of Assumption

This assumption can be argued from a rational agents perspective. Simply note that agents maximize profits under imperfect information, that is choose the optimal action that solves $a_j^* = \mathrm{argmax}_{a \in A_j} \{ \pi_j(a_j, a_{-j}; \theta) + \varepsilon(a_j) \}$, subject to an error term ε containing unobserved information. The conditional probability above is obtained, if that error term follows a Type I extreme value distribution (also termed Gumbel distribution). In this rational agent setting, the scaling parameter α can be interpreted as the relative influence of the observed over unobserved utility components. For small (large) α the conditional distribution becomes flat (peaked) because unobserved shocks dominate (are negligible).

Our real motivation however is different. First, taking an evolutionary perspective, we argue that agents that do not strictly maximise profits survive under moderate selection pressures. This entails also a behavioural perspective, i.e. real people have been observed to satisfice rather than maximize. Our second argument relates to the fact that our model of the economy is inherently incomplete with respect to reality. It describes a knowledge economy without frictions or market failures, not included in the model.

From that perspective, the scaling parameter α regulates the extent to which model mechanisms translate into realisations of the economy. As α becomes smaller the agents' choice distribution becomes flatter, and with that also the distribution over realisations of the whole economy.

This is particularly interesting in view of the empirical analysis. If the distribution is uni-modal and highly peaked (some highly likely outcome) it makes sense to consider maximum likelihood (i.e. assume that the realization is the most likely under the model). If the distribution is flat, why should we believe that we are at a maximum? A particular question, relevant to policy implications but yet to answer, is whether α affects the bi-modality (or multi-modality) of the system's distribution (i.e. the length of staying in a certain region).

Simulations

A Markov chain Monte Carlo (MCMC) algorithm samples from the space of actions A according to a stable distribution of the system that is implied by individual agent's conditional probabilities $p(a_j \mid a_{-j}, \theta)$ of Equation (5).

We design a Markov chain Monte Carlo (MCMC) sampler to draw from a distribution over joint agent actions $a = (a_1, \ldots, a_n)$, where each agent $j \in \{1, \ldots, n\}$ selects an action a_j to maximize its own profit $\pi_j(a_j, a_{-j}; \theta)$. The agents are situated on a directed graph, and their actions include both selecting a supplier and choosing an individual characteristic (i.e. price offered and R&D investment).

The MCMC algorithm constructs a reversible Markov chain over the joint action space by sequentially updating agents or cliques (subsets of agents), guided by the following principles:

1. Target Distribution: The conditional distribution for agent j is defined as a logit-response function:

$$p(a_j \mid a_{-j},\bm{\theta}) \propto \exp\left(\alpha \pi_j(a_j, a_{-j}; \bm{\theta}) \right)$$

where $\alpha > 0$ is an inverse temperature parameter. This defines the stationary distribution of the Markov chain as proportional to the product of these conditionals.

2. Gradient-Informed Proposal: Proposals for new actions $a_j' \sim q_j(a_j' \mid a_j, a_{-j})$ are generated using local numerical gradients of π_j with respect to a_j, biasing proposals toward locally improving directions. Reverse transition probabilities $q_j(a_j \mid a_j', a_{-j})$ are calculated accordingly.

3. Metropolis-Hastings Acceptance: Proposed actions are accepted with the usual Metropolis-Hastings probability:

$$r_j = \min\left\{1,\frac{p(a_j' \mid a_{-j},\bm{\theta})\, q_j(a_j \mid a_j', a_{-j})}{p(a_j \mid a_{-j},\bm{\theta})\, q_j(a_j' \mid a_j, a_{-j})} \right\}$$

ensuring detailed balance with respect to the logit-based target distribution.

4. Clique-Based Joint Updates: To improve mixing and capture strategic complementarities, the sampler occasionally performs joint updates of a clique $C \subset \{1, \ldots, n\}$. A joint proposal $a_C' \sim q_C(\cdot \mid a_C, a_{-C})$ is generated using the joint gradient of the combined clique profit:

$$\sum_{j \in C} \pi_j(a_j, a_{-j}; \bm{\theta})$$

$$\mathbf{a}_{\mathcal{C}}' = \mathbf{a}_{\mathcal{C}} + \eta \nabla_{\mathbf{a}_{\mathcal{C}}} \sum_{j \in \mathcal{C}} \pi_j(\mathbf{a}; \bm{\theta}) + \sigma \boldsymbol{\xi},\quad \boldsymbol{\xi} \sim \mathcal{N}(0,I)$$

$$\frac{p(\mathbf{a}')}{p(\mathbf{a})} = \prod_{j = 1}^n \frac{p(a_j' \mid a_1',\dots, a_{j-1}', a_{j+1},\dots, a_n)}{p(a_j \mid a_1',\dots, a_{j-1}', a_{j+1},\dots, a_n)}$$

$$\frac{p(\mathbf{a}')}{p(\mathbf{a})} = \prod_{j \in \mathcal{C}} \frac{p(a_j' \mid a_1',\dots, a_{j-1}', a_{j+1},\dots, a_n)}{p(a_j \mid a_1',\dots, a_{j-1}', a_{j+1},\dots, a_n)}$$

$$p(a_j' \mid a_{\mathcal{C},-j}', a_{\mathcal{C}^c}) \propto \exp\left(\alpha \pi_j(a_j', a_{\mathcal{C},-j}', a_{\mathcal{C}^c})\right)$$

$$\frac{p(\mathbf{a}')}{p(\mathbf{a})} = \exp\left(\alpha \sum_{j \in \mathcal{C}} [\pi_j(a_j', a_{\mathcal{C},-j}', a_{\mathcal{C}^c}) - \pi_j(a_j, a_{\mathcal{C},-j}', a_{\mathcal{C}^c})] \right)$$

$$r = \min\left\{1, \frac{p(\mathbf{a}')}{p(\mathbf{a})} \frac{q(\mathbf{a}_{\mathcal{C}} \mid \mathbf{a}_{\mathcal{C}}')}{q(\mathbf{a}_{\mathcal{C}}' \mid \mathbf{a}_{\mathcal{C}})} \right\}$$

5. Update Schedule: At each iteration:
   • With probability p_clique, a random clique is selected and jointly updated.
   • Otherwise, a random agent is selected and updated individually.

This hybrid approach allows the algorithm to efficiently explore the state space by leveraging local profit gradients while maintaining detailed balance and convergence to the logit-based stationary distribution.

3. Research results

This section presents our main observations on simulation experiments.

The output of simulating a model subject to some parameters in a MCMC is a collection of states, drawn from a stable state distribution implied by the model. The model is denoted Θ, the parameters that fully specify the model.

A statistic of interest can be calculated on each simulated state, say , $t(s)$ which then also follows a distribution that is implied by the state distribution. The average of the statistic over simulated states corresponds to the expected value of that statistic.
Changes in parameters result in changes of the distribution of statistics. Comparing distributions, or sample averages and sample variance, thus allows us to see aggregate effects of parameter changes.

Policy experiments can be considered as a change of the model or parameters. We implement global policies (tax credits on R&D as a reduction in cost of R&D), industrial policies, and individual policies (subsidies to individual firms). Any policy is financed by a reduction of the consumption budget of the consumers. Note that R&D spending does not lead to further processes such as hiring labor.

Assume R&D would be spent on labor (researchers' wages) — then R&D spending also adds to the income
and hence policy would be costless — it is only a redistribution of the income.

In the model, agents within the same industry are ex ante symmetric in the sense that their labels are arbitrary.

Policies that affect equilibrium states are considered here.
In fact, the MCMC simulation does not trace behavioral interdependence among agents,

such as the marginal effect of i’s R&D on i’s R&D. But marginal effects may not even be of interest.

4. Conclusion and Comments

Our concluding remarks are twofold. First, when network and R&D decisions are jointly determined, the economy tends toward a unique equilibrium under given structural conditions. The introduction of stochastic choice through the logit-based framework allows for bounded rationality and heterogeneity in firm behavior, producing richer and more realistic dynamics than a purely deterministic setting. Second, the MCMC simulations reveal how equilibrium realizations evolve as model parameters—such as selection intensity or innovation costs—change.

From a policy perspective, these findings imply that innovation support mechanisms should account for the heterogeneous structure of industrial linkages. Our study reveals that the collective outcome of individual R&D strategies depends crucially on who interacts with whom—and how strongly.

5. Outputs (Publications, presentations, etc.)

   The current project is close to a working paper and an upcoming presentation at seminer series of BETA (Bureau d'Economie Théorique et Appliquée) under the CSI (Creativity, Science and Innovation) axis.

Note: The accrued rate of the proposed study should be reflected to the report.

RESEARCHER NAME	SIGNATURE	DATE
Cemil Faruk Durmaz		26.10.2025

DIRECTOR NAME	SIGNATURE	DATE
Moritz Mueller		26.10.2025