A branching and recombination model of innovation

Koen Frenken[a] Luis R. Izquierdo[b] Paolo Zeppini[c]

December 2010

[a] *School of Innovation Sciences, Eindhoven University of Technology, k.frenken@tue.nl*

[b] *University of Burgos, luis@izquierdo.name*

[c] *CeNDEF, Faculty of Economics and Business, University of Amsterdam, p.zeppini@uva.nl*

Key words: agent based models, network externalities, technological transitions.

JEL classification: C15, O33.

Abstract

To explain the dynamics of technological transitions, we develop an agent based model based on network externalities and two different types of innovations. Recombinatorial innovations create short-cuts which speed up technological progress, allowing transitions that are impossible with only branching innovations. Our model replicates some stylized facts of technological transitions, such as punctuated equilibria, path dependency and technological lock-in. We find analytically a critical mass of innovators for successful innovations and technological transitions. Recombinant innovation counters network externalities, and calls for technological diversity as a key feature of technological transitions. An extensive simulation experiment shows that stronger network externalities are responsible for S-shaped utility and technological quality curves, indicating that a threshold of innovation probability is necessary to boost innovation. We finally introduce a policy view and interpret the innovation probability as the effort to foster technological change. A welfare measure including innovation costs presents an optimal interior value of innovation effort. The optimal innovation effort is strongly correlated with the number of recombinations, which further indicates how recombinant innovation is important in achieving a sustained technological progress at relatively low costs.

1 Introduction

Among the most challenging questions in the social sciences is the question how one can explain societal transitions. Transitions range from transitions in norms, in opinions, in preferences, and in technology use. It is the latter case we will refer to in the following though we feel that some elements of the model developed below may be more generally applicable.

We characterise transitions as large-scale changes that occur suddenly yet endogenously. This means that the time-scale at which a transition takes place in a particular context is considerably smaller than the time-scale at which transitions do not take place in that context. It also implies that we do not invoke an external cause (shock) to explain transitions.

To explain the dynamics of technological transitions, we develop a model where agents enjoy positive network externalities from using the same technology, while some agents, called entrepreneurs, ignore these externalities and introduce new technologies. Models of path dependence (David, 1985; Arthur, 1989) only explains how a technology becomes dominant in a population, and do not explain the emergence of new technological paths. The call for models that combine path creation and path dependence is legitimate, as they are fundamental aspects of transitions to sustainable technologies.

We assume that technologies form a graph (Vega-Redondo, 1994; Carayol and Dalle, 2007) which is evolving because entrepreneurs create new nodes. Remaining agents make decisions about technology and only adopt a new technology if it gives higher returns net of the switching costs. A specific feature in our model, under-explored so far, holds that technologies can be recombined (van den Bergh, 2008; van den Bergh and Zeppini-Rossi, 2008). Re-combinatorial innovations create short-cuts which speed up technological progress, allowing transitions that are impossible otherwise. Recombination short-cuts operate kind of re-wiring in the graph similar to "small world" networks (Watts and Strogatz, 1998), but in a dynamic environment.

Our model replicates some stylized facts of technological transitions, such as punctuated equilibria, path dependency and technological lock-in. A theoretical analysis shows that entrepreneurs can always break a technological lock-in, though the time required rises exponentially with the population size (Bruckner et al., 1996). We find analytically a critical mass of innovators for successful innovations and technological transitions. The model also account for potentially good innovations that became unsuccessful only because of network externalities. Recombinant innovation counters network externalities, and calls for technological diversity as a key feature of technological transitions.

An extensive simulation experiment gives the following results. When the population size is small and network externalities are weak, the mean utility and minimum quality of technologies are increasing in the probability of innovation, with a saturation effect at large values. Stronger network externalities make utility and quality curves S-shaped, indicating that a threshold of innovation probability is necessary to boost innovation. When strong externalities are accompanied by a large population size, a local minimum of utility appears at low values of the probability of innovation, indicating an initial detrimental effect of innovation.

We finally introduce a policy view and interpret the innovation probability as the effort to foster technological change. A welfare measure including innovation costs presents an optimal interior value of innovation effort, which follows from the S-shape of the utility and indicates that neither too low or too high efforts are advisable for innovation policy. The optimal innovation effort is strongly correlated with the number of recombinations, which further indicates how recombinant innovation is important in achieving a sustained technological progress at relatively low costs.

The paper is organised as follows. Section 2 presents the model and contains an analytic study of its properties. Section 3 contains some simulation examples and the numerical analysis of an extensive simulation experiment. Section 4 concludes, also indicating the direction for possible extensions of the model.

2 Our model

Here we present formally our agent based model of technological innovations. For a more detailed description of the model refer to appendix A. Let there be a population of N agents ($N \geq 2$), which in every period face a decision about which technology to adopt. Technologies belong to an expanding set, which is built by agents themselves through an innovation process. Given the technology set A_t of period t, agents decide based on a utility $u_{\alpha,t}$, where $\alpha \in A_t$

indicates the technology adopted. The utility from using technology α comes from an intrinsic quality l_α and from the positive externalities that other users of α exercise on the single agent:

$$u_{\alpha,t} = l_\alpha + e n_{\alpha,t} \tag{1}$$

where the parameter $e \in [0,1]$ measures the strength of network externalities, while $n_{\alpha,t}$ indicates the number of agents using α in period t. Technologies form a directed graph of which they represent the nodes, while the links express the generational relation. The graph evolves, due to the action of agents, which is either a technology adoption or an innovation. Adoption decisions determine the population size of each technology, or node in the graph. Innovation by an agent generates a new technology, and then a new node and at least one new link. We can assume without any loss of generality that in the beginning only one technology is available. If only branching is possible, this is obvious, because each starting technology evolves independently. If recombinant innovation is possible, the initial condition where more technology are present can always be replicated by starting with only one technology, if one can wait long enough. This has to do with the ergodicity of the model with recombinant innovation, an issue that we will address later in the paper.

2.1 Adoption decision

In every period t an agent may be drawn as innovator with some probability p. If an agent is not innovator it evaluates and compare the utility from adopting each available technology in the set A_t. All non-innovator agents decide synchronously which technology to adopt. The decision is actually about whether to stay with the actual technology or to switch to the more attractive among other technologies. Such decision involves a third factor, the switching cost, which we assume equal to the geodesic technological distance between the used and the new technology. For instance, switching from technology α to technology β takes place as soon as the following condition realizes:

$$u_{\beta,t} - d_{\alpha\beta} > u_{\alpha,t} \tag{2}$$

Let all technologies be part of a connected graph with the technological distance between α and β given by the geodesic distance $d_{\alpha\beta}$ (with $d_{\alpha\alpha} = 0$). We assume that this quantity also represents the switching costs from one to the other technology. If the difference $\Delta u_{\alpha\beta} = u_{\beta,t} - u_{\alpha,t} - d_{\alpha\beta}$ is positive, agents will migrate from α to β. in case of $\Delta u_{\alpha\beta} = 0$, the old technology is mantained. Since more than two technologies are present in the network in general, agents search for the best one. If two technologies β and γ present the same benefits from switching, that is if $\Delta u_{\alpha\beta} = \Delta u_{\alpha\gamma}$, a random decision is taken. In the following we assume that switching costs for adjacent technologies are constant in time and the same for all technologies.

2.2 Innovation

Technological opportunities change through time, innovation being possible. In each period any agent can innovate with probability p, introducing a new technology that represents a quality improvement with respect to the technology previously used. For simplicity we assume that quality improvements are always equal to one. There are two channels of innovation: *branching* and *recombination*. In the first case, one or more agents from the same technology innovate creating a new technology that “branches” from the old one. In the second case, at least two agents from two different technologies join to create the recombinant innovation. In the technology graph a recombinant technology has at least two incoming links from different parent technologies, while with branching the incoming link is always one (figure 1).

In the case of branching the improvement is a unitary step up over the parent technology. If β is an innovation that branches from technology α, we have:

$$l_\beta = l_\alpha + 1 \qquad \textit{branching} \tag{3}$$

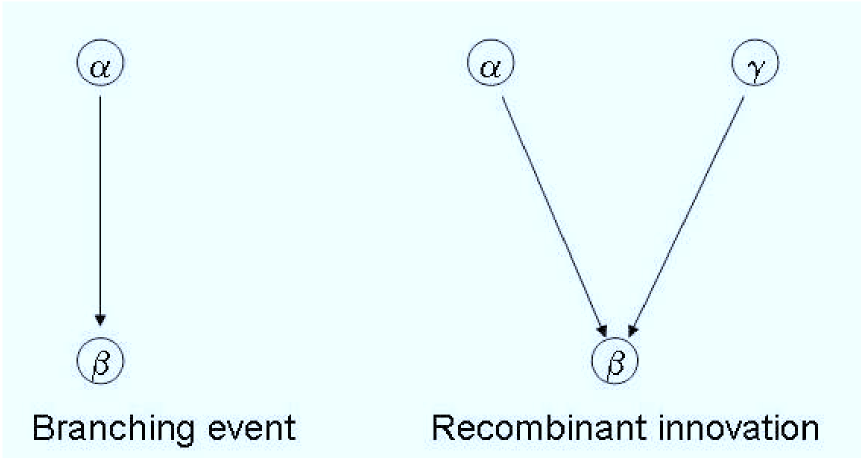

Figure 1: Intuitive representation of innovation events in one period: branching (left) and recombinant innovation (right).

When recombinant innovation arises, the quality of the innovation is assumed to be a unit higher than the maximum quality of recombinant technologies. If α and γ recombine to give the innovative technology β, we write:

$$l_\beta = \max\{l_\alpha, l_\gamma, \dots\} + 1 \qquad \textit{recombination} \tag{4}$$

In general, if m technologies recombine the quality of the innovation will be one unit higher than the quality of the best among these m technologies[1]. By assumption the innovator agents stick to their new born technology at least for one period. The model with or without recombinant innovation present some sharp difference, which suggest to address them separately. This is what we do in the next two sections.

2.3 Branching innovations

This specification of the model is more suited to describe linear technological progress. Technological quality improvement are always unitary for all technologies. Innovations can only have one parent technology and the technological graph appears as a star, with a number of branches that reaches the total size of the population in the long run. When this attracting state is reached, the technological variety does not change anymore. In other words, the model evolves towards a situation where alternative technologies branch out, each one being carried on by a single agent. Figure 2 reports an example of a simulation run of the model[2]. In this example we have 5 agents, we set $e = 0.1$ (externalities) and $p = 0.1$ (probability of innovation), and we let the model run for $T = 50$ periods. The graphic representation of the simulated model shown by figure 2 contains information about the number of technologies invented (we start we

[1]The above assumption is critical and debatable: it can also be possible that parent technologies with lower quality act as "bottlenecks" and keep low the quality of the recombinant innovation. This issue raises a more fundamental question about the nature of recombinant innovation. Parent technologies should have a certain degree of complementarity in order to recombine. If the parent technologies become the modules of the recombinant innovation as a complex system, the quality of this new technology is likely to be a sort of average of the quality levels of parent technologies. The difficulty of finding a suitable specification for the dynamics of technological progress of recombinant innovation also comes from the low dimensionality of quality: using one scalar parameter to describe performance is probably oversimplifying. It applies quite well to incremental innovation and to the schematization of branching, where the innovative technologies performs the same task of the old ones. It is maybe too restrictive for recombinant innovation, where the different parents are not necessarily substitute technologies, and then the idea of quality improvement needs to be extended.

[2]The model has been implemented in NetLogo

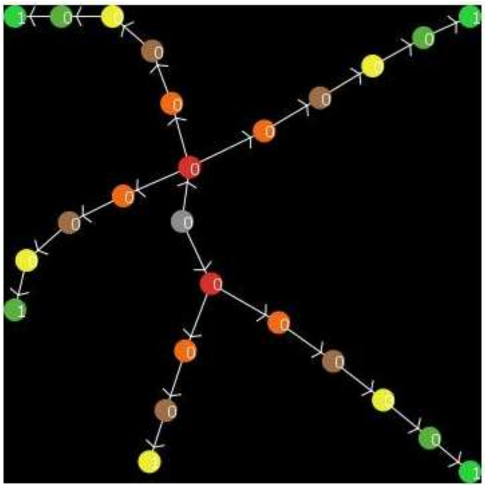

Figure 2: Example of simulation of the model with only branching. Here we have $N = 5$ agents, $e = 0.1$, $p = 0.1$ and the time horizon is $T = 50$. The numbers in the balls represent the population size of that technology.

just one technology), how many agents are in each technology (numeric label of the balls) and the quality level of each technology (colour of the balls). Moreover the representation show the directed links that connect technologies and that tell the genealogy of the final configuration. In the example of figure 2 we have five branching events, three stemming from one technology and two from another one. One agent inhabits the latest technology of each branch. In three cases the branch reaches the maximum quality of the system, namely level six (the initial technology has level zero). Another branch stops at level five, while the last one only reaches level four.

The system of technologies where branching is the only channel of innovation is non-ergodic. Once branches are formed, each with one agent, the total utility (or welfare) increases in a deterministic fashion, at a rate that depends on p.

In the following we derive some analytical properties for the branching model. We start with asking the following question: for what population size of the parent technology is an innovation always successful? With successful innovation we mean an innovation that is capable of attracting another agent and do not loses its innovator. Assume we are in the initial stage, with only technology α, and assume only one agent innovates producing technology β . The condition for successful innovation is that utility from switching to it is larger than utility from staying in the old technology, net of switching costs $d_{\alpha\beta}$ that we normalize to one:

$$l_\alpha + 1 + en_{\beta,t=1} - 1 > l_\alpha + en_{\alpha,t=1}$$

In the left hand side the increase in quality and the switching cost offset each other. Considering that we are in the initial stage and only one agent is innovator, we have $n_{\beta,t=1} = 1$ and $n_{\alpha,t=1} = N - 1$, which gives the condition

$$N < 2 \tag{5}$$

If we are considering a generic time t, and α is not the initial technology, we assume that M agents are occupying other technologies, so that $n_{\alpha,t=1} = N - M - 1$, and the condition for successful innovation becomes

$$N < 2 + M \tag{6}$$

One important remark: agents are myopic, in that they do not consider their own positive contribution to the utility. Moreover, they are not strategic, missing to take into account the contribution of other agents' actions. Based on these assumptions, when we write conditions for agents action we always assume the point of view of one single agent that computes utility as if all other agents do not switch.

Conditions (5) and (6) tell that a larger population of agents makes successful innovation more difficult, due to network externalities. Condition (6) tells also that a more diversified system gives a higher probability of successful innovations, because each technology is less populated. These two conditions are quite stringent. But by no means they rule out the possibility of successful branching, in the sense that it is still possible to escape the starting technology (escape lock-in) even when they are not met. What we need is two successive innovations. This consideration introduces the following question: what geodesic distance needs to be covered with successive innovations within a branch, in order to have a successful innovation? In other words, how many successive innovations make a successful branching event, where an agent (we consider a single innovator) leaves the original technology permanently? We address such question writing down the condition for which it is not profitable for the innovator to go back from its last invented technology β (the frontier technology) to the original technology α:

$$u_\beta \geq u_\alpha - d^*_{\alpha\beta} \tag{7}$$

where $d^*_{\alpha\beta}$ expresses the total technological distance from α to β. If technological recombination is impossible, this is equal to the difference in qualities levels $d^*_{\alpha\beta} = l_\beta - l_\alpha$. With only one innovator in β and a generic number $n_{\alpha,t}$ of agents in α, inequality (7) becomes

$$l_\beta + e \geq l_\alpha + en_{\alpha,t} - (l_\beta - l_\alpha) \tag{8}$$

which can be rearranged in the following condition:

$$l_\beta - l_\alpha \geq \frac{en_{\alpha,t} - 1}{2} \tag{9}$$

Condition (9) says that stronger network externalities and larger population sizes ask for bigger improvements in technological quality, in order for branching innovation events to be successful. If we have a generic number of innovators $m_{\beta,t}$, the condition can be easily generalized by substituting $en_{\alpha,t} - 1$ with $e(n_{\alpha,t} - m_{\beta,t})$ in the numerator of (9).

We now ask a slightly different question: how many agents have to "co-invent" for the others to follow suit? Let consider a generic time t, with n agents in technology α, and assume that m agents co-invent technology β in that period. Assume that no other innovation events characterize technologies α and β in periods t and $t+1$. The m inventors stick to the innovation in period $t+1$ by assumption. Two cases are possible: 1a) also the other $n-m$ agents follow suit and adopt technology β, or 1b) the other $n-m$ agents stay with the old technology α and the m innovators can do one of the two following actions: 2a) they decide to remain with technology β, or 2b) they go back to technology α. Case (1a) represents a *transition*, with an increase in the minimum quality level of all used technologies. If this case occurs in period $t+1$, it lasts also in the following period, and only another innovation event can change the population of technology β. Case (2a) represents a *branching* event, of which the transition event (1a) is a particular case. Case (2b) refers to an unsuccessful innovation, which we will refer to as *lock-in*, while cases (1a) and (2a) are successful. Figure 3 depicts the possible cases.

In what follows we derive mathematical conditions for cases (1a) and (2a). The case (2b) is complementary to (2a). The condition for a transition (case 1a) is the following:

$$l_\alpha + 1 + em - 1 \geq l_\alpha + e(n - m) \quad \Rightarrow \quad m \geq \frac{n}{2} \tag{10}$$

If more than half of the agents with technology α are selected to be innovators in period t, all other agents will follow, and a technological transition realizes. The threshold (10) only depends on the population size n, and not on the intensity of externalities e. This is because a transition requires that at least half the population of a technology co-invent, and if this is the

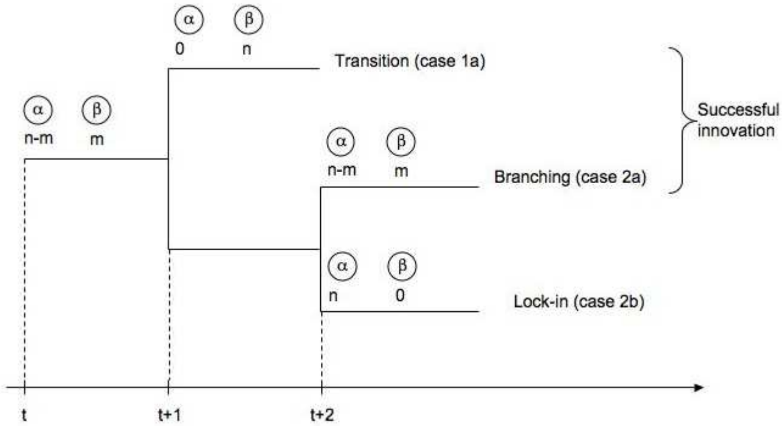

Figure 3: Possible cases of branching innovation by m co-innovators, starting from a technology α with n agents at a generic time t.

case, weak or strong externalities do not make a difference. If we have less innovators instead, a branching event is still possible, with case $(2a)$. The condition is:

$$l_\alpha + 1 + em \geq l_\alpha + e(n-m) - 1 \quad \Rightarrow \quad m \geq \frac{n}{2} - \frac{1}{e} \tag{11}$$

This condition is obtained by requiring that the innovators do not go back to the old technology α at time $t+2$. Figure 4 summarizes the different cases and relative threshold values for the number of co-innovators m. Because of condition (11), the stronger network externalities e are, the higher is the threshold level of co-innovator for escaping lock in. At the macro level the

Figure 4: Summary of conditions on the number m of co-innovators for a *transition*, and/or for a *branching* event, starting from a technology α with n agents at a generic time t.

properties outlined above determine the following features of the model. First of all, once a transition occurs, the minimum quality does not step back. Total utility can go down instead, because innovators give up in terms of network externalities. The non-negative (although not deterministic) rate of increase of quality is behind the non-negative rate of the total number of technologies in use. But by no means the entropy of the system always increase. Entropy may decrease in some case, if agents happen to reorganize in a way that more of them use the same technology. A last feature of the model with only branching is that agents never adopt an unused technology. This is not the case when recombinant innovation is possible.

2.4 Branching and recombinant innovations

With recombinant innovation being possible, two or more different technologies can be the parents of a new one if at least one agent from each of these technologies is drawn as an innovator. Branching is still possible, if only one agent is drawn as innovator in one period. By assumption, multiple innovators in a same period always recombine.[3] The connectivity of the resulting technology graph is higher than the case with only branching. Figure 5 reports

[3]An extension of the present model, called "Partitioning model" considers all possible cases with multiple innovators in a same technology.

an example of simulation run in the same condition as figure 2, but with recombination being possible. From this example we see already by visual inspection a first virtue of recombinant

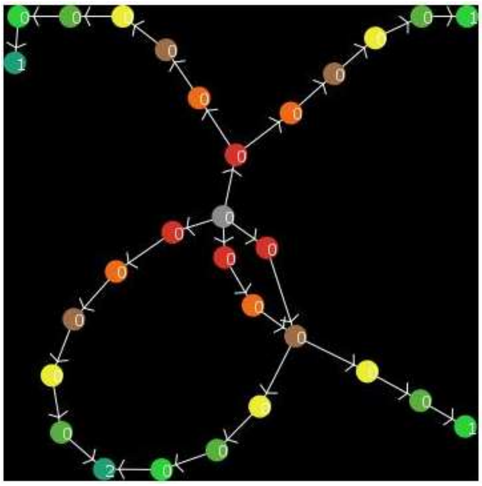

Figure 5: Example of simulation of the model with branching and recombinant innovation. Here we have $N = 5$ agents, $e = 0.1$, $p = 0.1$ and the time horizon is $T = 50$. The numbers in the balls represent the population size of that technology.

innovation, which is the possibility of "short-cuts" to higher quality technologies: agents that recombine will profit from the other parent technology and make a jump in quality larger than one. This is the case of one of the two agents that inhabit one of the two highest level technology in the example: such agent had to make only six instead of seven steps to get there. The short-cut property translates into a higher rate of increase in minimum quality of the overall technology graph, and consequently in a higher probability of technological transitions.

There is an important and fundamental difference between the model with and without recombinant innovation. We have seen that if branching is the only channel of innovation, the technology graph reaches an attractor represented by a state with as many branches as agents, and the system is non-ergodic. The system with recombinant innovation is ergodic instead: with positive probability is reaches the state where all agents use the same technology. The quality of such technology is higher than the starting technology, but the state variable is the same (technology adoption shares, where one technology has share equal to one). We will cal this state the *recurring state*. This state will be reached surely in a finite time. After this time, everything goes as if one starts the model again, only with a different initial quality. Only quality tells that there was a progress, but the recurring state at different times is indistinguishable.

The probability of recombinant innovation is the joint probability that two or more agents are drawn to be innovators minus the probability that these agents are using the same technology. Disregarding the negative contribution of this second term, we may affirm that the probability of recombinant innovation of m technologies is proportional to p^m. This tells that recombinant innovations that involve more technologies are less likely.

The probability of having a recombinant innovation with any number of parent technologies is positively affected by the following factors: the number of agents n, the number of used technologies, the entropy of the system. The positive effect of entropy is due to the fact that more balanced distributions of agents across technologies give a higher probability of recombination. In other words, the marginal effect of taking an agent from a crowded technology and giving him a new technology is positive (Van den Bergh and Zeppini-Rossi 2008).

3 Numerical analysis

3.1 Simulation examples

Our model is actually a model of economic growth, which includes technological transitions. An implicit aspect of technological transitions is the positive rate of change in quality and mean utility of agents. With *technological transtition* we mean a positive increase in the minimum quality level of used technologies. A central role in our model is the probability of innovation p. Network externalities e and agents population size N may be seen as given parameters that characterize the social system. But p is a parameter open to policy intervention, through government investments aimed at fostering research and development, for instance. In order to give a flavour of the role played by the probability of innovation p in the model we run four simulation examples for four different values of this parameter. Il all these examples we consider $N = 50$ agents, with a magnitude of network externalities equal to $e = 0.5$. We run the simulations for 50 periods. In the first example we set $p = 0.1$ (figure 6): the visual

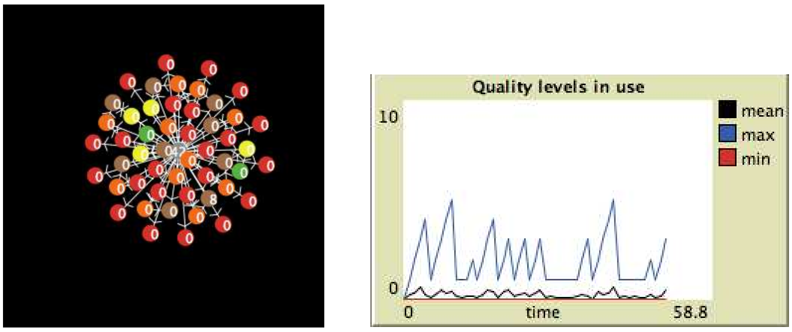

Figure 6: Simulation run with $p = 0.1$ ($N = 50$, $e = 0.5$ and $T = 50$ periods). Left: technology graph. Right: Minimum quality level of used technologies (red line), maximum level (blue line) and mean level (black line).

representation of the technology graph shows that 42 agents remain in the initial technology and 8 agents are in a technology with higher quality. Actually this configuration gives an illusion, because these agents are going to go back in the lower level technology in next periods (unless they are again drawn as innovators). This is because the population size is too large the network externalities too strong and the switching costs too low for innovative agents to remain in better technologies. This fact explains why the maximum level of used technologies presents many positive swings, while the minimum level stays at zero (right part of figure 6). We may say that in this setting we have no technological transitions and zero growth. The next example has $p = 0.2$. In this case we have three transitions, as the three steps of the minimum quality shows in the right part of figure 7. The three transitions reflect in the cluster structure of the technology graph in the left part of the figure. We have four clusters in the 50 periods of this simulation run: one builds around the initial technology. A second one stems from an innovation that was successful initially, but eventually was left by all agents. A third cluster is successful instead, in the sense that lately gives place to a new cluster, thanks to the fact that the innovative technology is able to attract all agents in the population. Here we see what we analyze theoretically in section 2.3: a critical size of followers is necessary to make an innovation successful. Increasing further the probability of innovation we obtain even more transitions with faster growth. Figure 8 reports an example with $p = 0.3$, while figure 9 has $p = 0.5$. In both examples the maximum level of quality increases continuously, indicating that in this setting $p = 0.3$ was enough to ensure a successful innovation in every time step. The variability of the system decreases substantially, and the model acquires more and more a deterministic character. The technology graph on its turn assumes the form of a chain, with

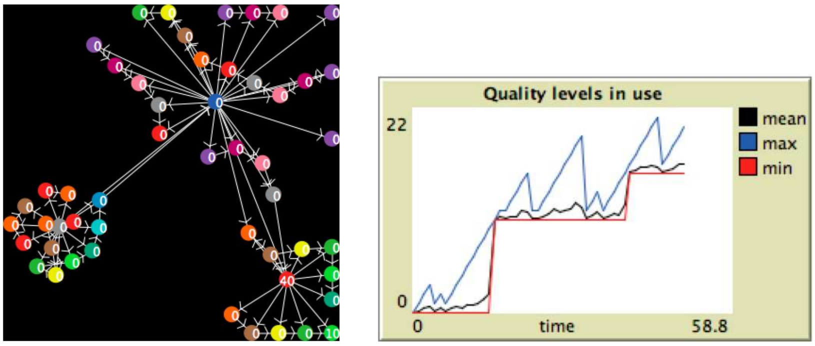

Figure 7: Simulation run with $p = 0.2$ ($N = 50$, $e = 0.5$ and $T = 50$ periods). Left: technology graph. Right: Minimum quality level of used technologies (red line), maximum level (blue line) and mean level (black line).

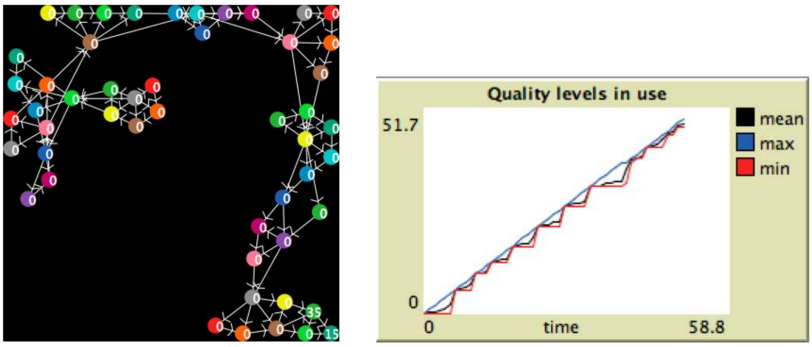

Figure 8: Simulation run with $p = 0.3$ ($N = 50$, $e = 0.5$ and $T = 50$ periods). Left: technology graph. Right: Minimum quality level of used technologies (red line), maximum level (blue line) and mean level (black line).

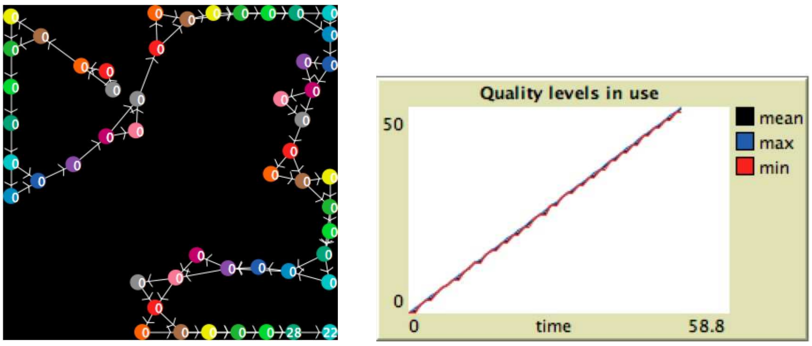

Figure 9: Simulation run with $p = 0.5$ ($N = 50$, $e = 0.5$ and $T = 50$ periods). Left: technology graph. Right: Minimum quality level of used technologies (red line), maximum level (blue line) and mean level (black line).

little space for clustering.

3.2 A simulation experiment

In this final section we report the analysis of a systematic simulation experiment, that aims at unveiling the effect of the innovation probability p on the model, in different conditions of agents population, namely its size N and the network decision externalities e. In such simulation experiment we considered a time horizon of 50 steps, and averaged results over 10 repetitions. We then analyze simulation results by looking at four variables: the minimum quality of used technologies, the mean utility of agents, the total number of recombinations and the accumulated entropy. the first two variables are computed at the end of the simulation run, while the second two are cumulative variables, being made of the contributions of all periods in the simulation time horizon. In particular, the entropy measures at the same time the variety and the symmetry of the system. In a given period t, the entropy of the system is defined as

$$E_t = - \sum_{\alpha \in A_t} \frac{n_{\alpha,t}}{N} \log_2 \frac{n_{\alpha,t}}{N} \tag{12}$$

Let us initially consider a situation with only 2 agents and very low decision externalities, $e = 0.1$. The left part of figure 10 reports the cumulation of two quantities over the time horizon considered (100 steps), namely the accumulated entropy and the total number of recombinations. On the right part of the figure, instead, we have the values of minimum quality of used technologies and mean utility of agents at the last step of the simulation run (the 100^{th}). Two main facts arise from a quick visual inspection of the graphs: first, there is one value of

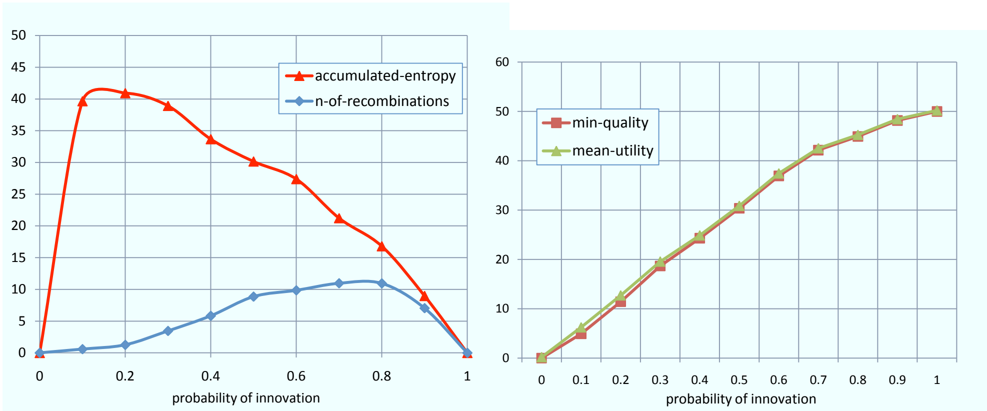

Figure 10: Simulation with $e = 0.1$ and $N = 2$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

the probability of innovation p that is an interior maximum for the accumulated entropy and a larger value of p that is a maximum for the number of recombinations. In other words, the value of p that maximizes entropy is much lower than the value that leads to a maximum number of recombinations. Secondly, the minimum quality of technologies and the mean utility of agents are strongly correlated, as expected considering how utility is defined (see equation 1), and they are strictly increasing in the probability p.

What is the effect of stronger network externalities? Figure 11 reports the results for $e = 0.5$. The effect is moderate, and any effect is hardly detected. Only a slightly lower maximum for

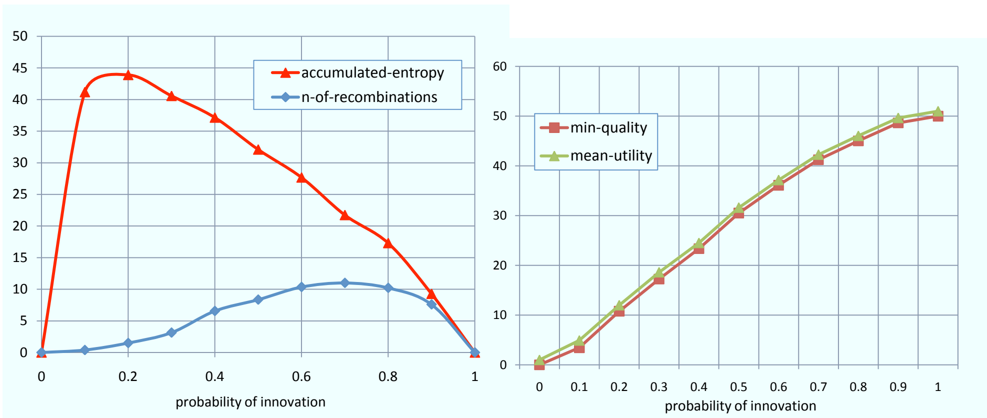

Figure 11: Simulation with $e = 0.5$ and $N = 2$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

the total number of recombinations. What is the effect of a few more agents, instead? Figure

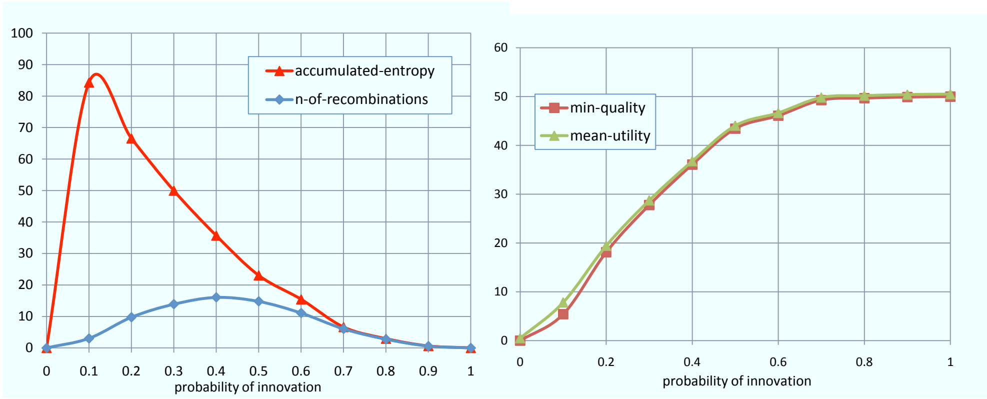

Figure 12: Simulation with $e = 0.1$ and $N = 5$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

12 considers a case with 5 agents and $e = 0.1$. Here the differences are more evident: first of all, the value of p that leads to a maximum number of recombinations is quite lower than with only two agents. Secondly, the minimum quality and mean utility reached at the end of the simulation time horizon seem to saturate for large values of p. In other words, with five agents when p is relatively large the marginal contribution to utility of increasing p further is negligible. This suggests that pushing p to large values over some saturation level is useless. Notice how the limit value of the utility for $p \to 1$ is predictable: it coincides with the time horizon of the simulation (50 in this case). This means that when agents innovate with probability one, quality increases at every time step and the model is fully deterministic.

The combined effect of a few more agent and stronger externalities is not big: with $N = 5$ and $e = 0.5$ we obtain a picture much similar to the case $N = 5$ and $e = 0.1$. This indicates that for small sizes of agents population, increasing the number of agents has a much stronger than increasing the magnitude of network externalities' in agents decisions.

Let us now consider 10 agents. Figures 13, 14 and 15 show the results for this case, with $e = 0.1$, $e = 0.5$ and $e = 1$, respectively. We focus on the graphs of accumulated quantities, first (left part of the figures). Increasing the magnitude of network externalities e, the maximum

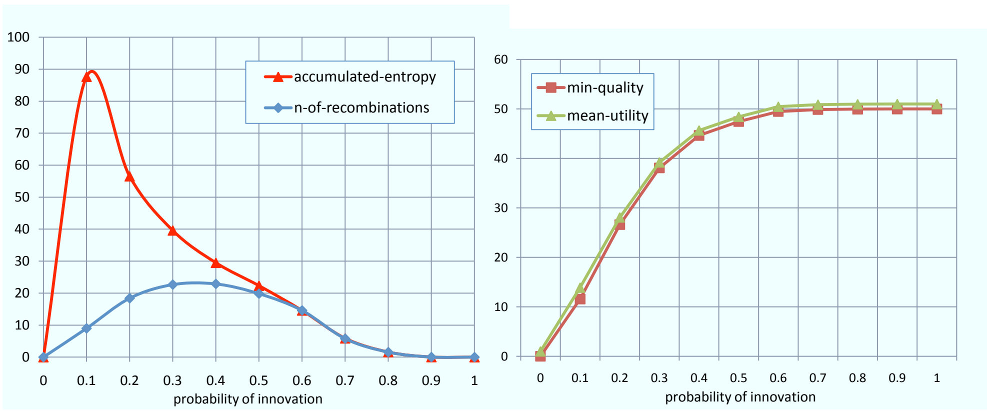

Figure 13: Simulation with $e = 0.1$ and $N = 10$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

of accumulated entropy goes down considerably, while such maximum is attained for slightly larger values of p. The curve of the total number of recombinations is almost the same for the three values of e, instead. We may say that stronger network externalities lead to lower entropy, increasing the level of coordination of agents in the system. Looking at the levels of quality and utility in the final period (right part of the figures) we see that 10 agents produce a stronger saturation effect of p: when $p > 0.6$, a further increase of this probability hardly gives any benefit, with no much differences for different values of e. Now with 10 agents we se clearly that in the limit $p \to 1$ the minimum quality level gets to a value equal to the time horizon of the simulation, and the utility gets to the same plus the externality of all agents, that is $T + eN$.

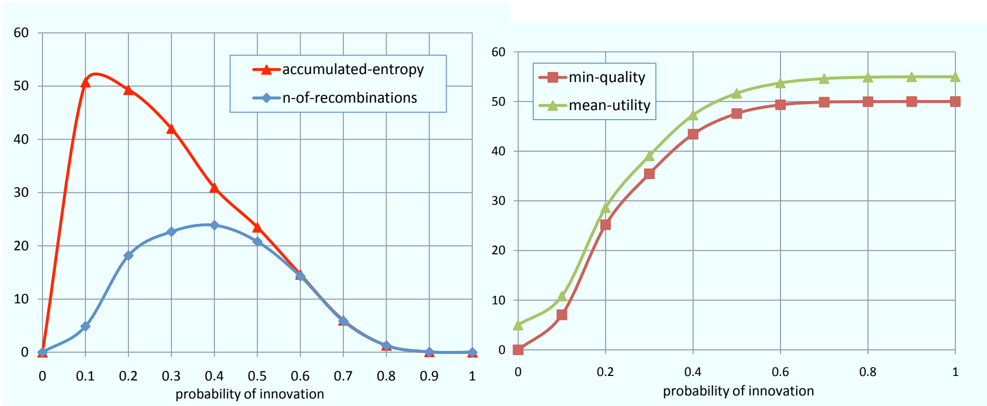

Figure 14: Simulation with $e = 0.5$ and $N = 10$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

The simulations with 10 agents present a second effect from network externalities: the occurrence of an S-shaped behaviour of the quality and utility curves, which become convex for low values of p. This means that when agents are more than a few and network externalities are intense, there is kind of a threshold value of the probability of innovation, after which technological transitions take off with consistent increases in quality and utility.

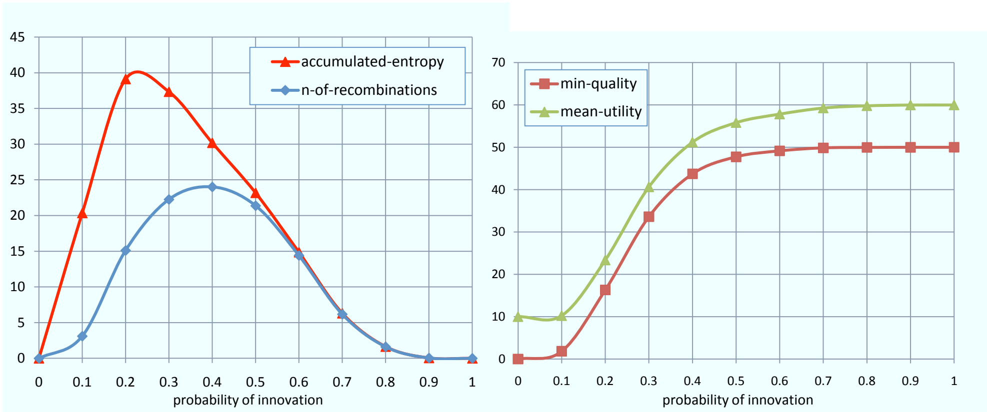

Figure 15: Simulation with $e = 1$ and $N = 10$. Left: accumulated entropy and total number of recombinations. Right: Minimum quality of used technologies and mean utility across agents.

The effects considered above become more evident with a larger population size. In the two figures that follow we consider the cases $N = 20$ (figure 16) and $N = 50$ (figure 17). One effect of increasing further N and e is the lower maximum point of total recombinations, which becomes almost coincident with the maximum point of the entropy. In other words, with many agents and strong network externalities, there is a value of the probability of innovation p that maximizes both entropy and number of recombinations. A second interesting effect is the non monotonicity of utility for low values of p. We have seen before that with 10 agents a strong magnitude of network externalities e makes the utility curve S-shaped. With more agents, this effect is more dramatic, producing an utility curve which is decreasing for low values of p and increasing for larger values. The same effect is not found for the quality level of technologies, because this is always 0 when $p = 0$. Beside the non-monotonicity of the utility curve, the combined effect of larger population size and network externalities deviates from the S-shape also for mid values of p: as the bottom part of figure 17 shows, the marginal effect of increasing p in the range $(0.3, 0.6)$ is quite irregular, suggesting a local $S-shape$ before getting to saturation.

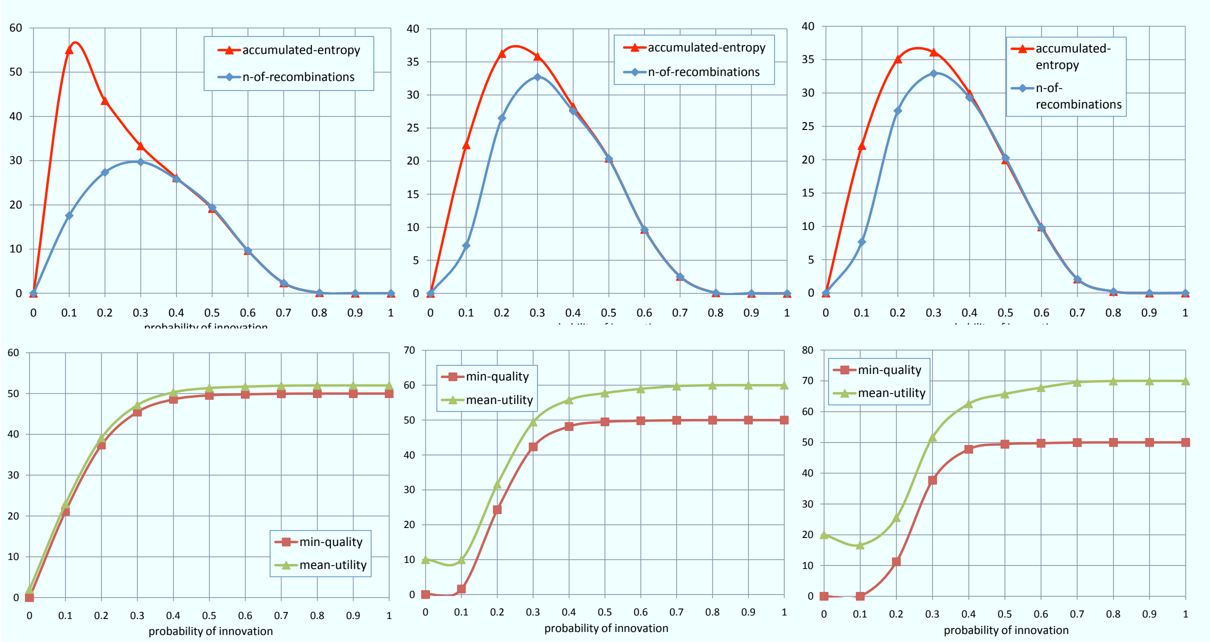

Figure 16: Simulation with $N = 20$. Top: accumulated entropy and total number of recombinations. Bottom: Minimum quality of used technologies and mean utility across agents. Left: $e = 0.1$. Centre: $e = 0.5$. Right: $e = 1$.

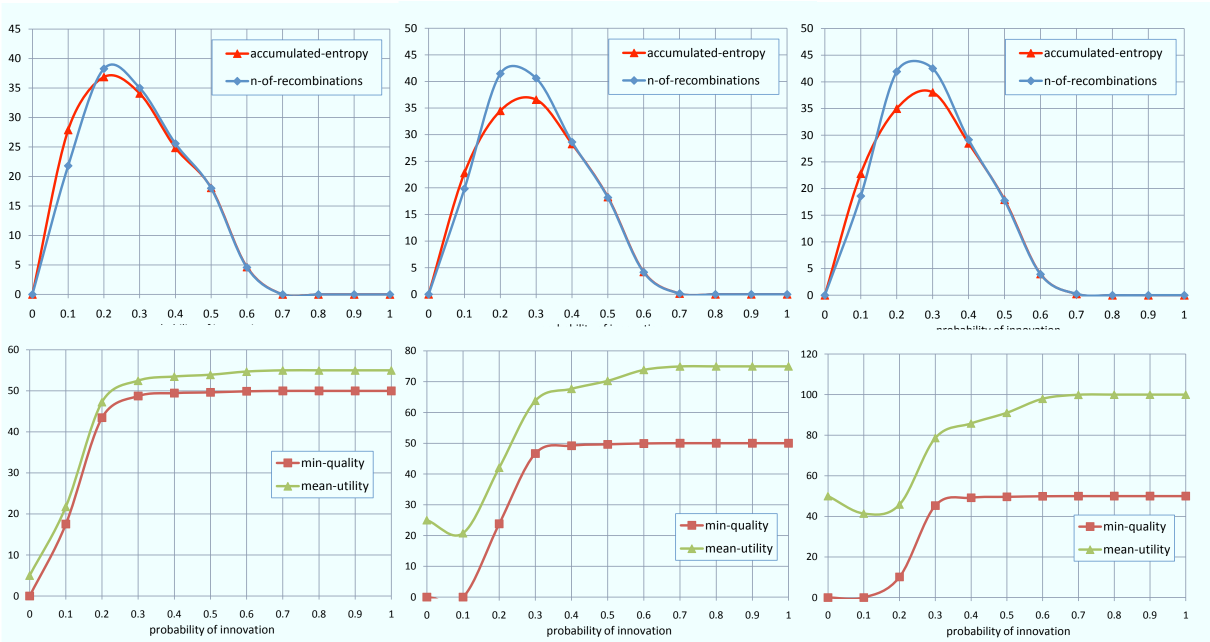

Figure 17: Simulation with $N = 50$. Top: accumulated entropy and total number of recombinations. Bottom: Minimum quality of used technologies and mean utility across agents. Left: $e = 0.1$. Centre: $e = 0.5$. Right: $e = 1$.

As a final piece of analysis we look at the welfare of the system, considering also the costs of innovation. A natural way of measuring innovation costs is to interpret the probability of innovation p as the effort devoted by a government policy to subsidize innovation. In a first approximation, the probability of innovation is assumed to be directly proportional to these costs. Then we measure the welfare as the mean utility net of such costs:

$$w =< u(p) >_N -cp \tag{13}$$

where the proportionality parameter c measure the efficiency of the innovation policy (the lower c, the more efficient the policy investment). Our main interest is to see whether there exist an optimal level of policy investment p that maximizes welfare. Figure 18 reports the value of w

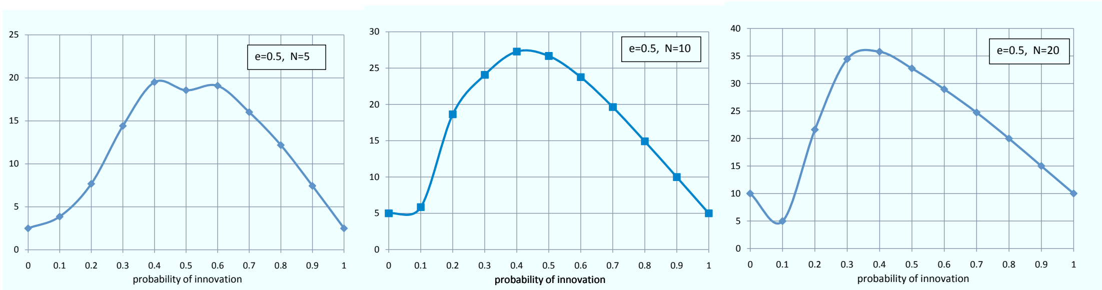

Figure 18: Measure of welfare w for different values of the innovation probability/cost. Left: $N = 5$. Centre: $N = 10$. Right: $N = 20$ ($e = 0.5$).

as a function of p for three different sizes of agents population ($N = 5$, $N = 10$ and $N = 20$), always with a fixed level of network externalities ($e = 0.5$). As we can see, there is clearly an interior solution for the welfare maximization problem (at least one). Moreover, as the population size increases, small values of p give a minimum of the welfare. This character of the dependence of welfare on p directly follows from the qualitative behaviour of mean utility that we have analyzed before: the S-shaped profile of utility lead to an interior maximum of welfare. This indicates that innovation policy efforts should be large enough in order to trigger the occurrence of innovations, but not too large because after some level the marginal benefits of any further investment is almost zero.

4 Conclusions

In this paper we have proposed an agent based model of technological change, that puts emphasis on two mechanisms of innovation, namely technological branching and technological recombination. The action of innovating agents is central in the model, which is an aspect that recognize the important role of entrepreneurs in technological change. Innovation is made by innovators but it is shaped by adopters. The model accounts for the stylize facts of technological changes such as increasing returns on adoption and path dependence, introducing a positive network externality in the utility of agents.

The model is able to replicate qualitatively some stylized facts of technological transitions: innovations occur in a punctuated and irregular fashion. Moreover the model captures the path dependency of technological change as well as its extreme effect leading to technological lock-in. The model also is able to account for the occurrence of potentially good innovations that became unsuccessful only because of network externalities.

We have derived a number of analytical and numerical results about this model. We have shown analytically that innovation size matters, and a critical mass of innovators is necessary for an innovation to be successful and determine a technological transition by attracting all the agents of the parent technology. In the case where only branching innovation can take place (recombination being excluded) we have derived the analytical condition for successful innovations, which links the agent population size to the magnitude of network externalities: for a given population size, the weaker network externalities are, the more often technological transitions occur.

By running an extensive simulation experiment we have analyzed the role of the probability of innovation in different conditions of population size and network externalities, obtaining the following results: there are interior solutions for the maximization problems of accumulated entropy of the system and for the total number of recombinations in terms of the innovation probability. These interior solutions do not coincide. When the population size is small and network externalities are weak, the accumulated entropy is maximum for a value of the probability of innovation smaller than 0.5, while the total number of recombinations is maximum for a value larger than 0.5. By increasing population size and network externalities these interior solutions get closer to each other, until they almost coincide for $N = 50$ and $e = 1$. The mean utility of agents and the minimum quality of used technologies are strongly linked in our model. When the population size is small and network externalities weak, these two quantities are increasing in the probability of innovation, with a behaviour showing a saturation effect for large values of this probability. In other words, when the probability of innovation is above some level, the marginal effect of any further increase is negligible. A stronger network externalities effect makes the utility and quality level curves S-shaped, indicating that a threshold value of innovation probability is necessary in order to boost innovation. When such strong network externalities are accompanied by a large population size, an interior minimum of utility appears: if we increase the probability of innovation starting from zero, initially the mean utility goes down, indicating a detrimental effect of innovation. Above such interior minimum, the utility returns to be increasing in the innovation probability, although it shows different values of the rate of increase before getting to saturation.

Finally we introduced a policy view in the analysis, by interpreting the innovation probability as the effort made in order to foster technological change. By looking at a welfare measure that takes into account the costs of such a policy, beside the utility of agents, we find the existence of at least one optimal interior value of the innovation probability. The existence of such an interior optimum comes from the S-shaped behaviour of the utility. Moreover we have found evidence of possible multiple interior solutions for the welfare maximization problem. This is due to the fact that the utility deviates from the S-shaped behaviour not only for low values of the innovation probability, but also for mid values, when both the population size and the network externalities are large.

Appendix A Model description

In the model there is a constant population of N agents ($N \geq 1$) and an evolving set of technologies that agents may use. The model runs in discrete time-steps. At any given time-step t, each agent is using one and only one technology, and derives a certain utility from doing so. The utility $u_{\alpha,t}$ gained at time-step t when using technology α comes from an intrinsic quality of the technology l_α and from the positive externalities that other users of α exercise (equation 1 in the main text):

$$u_{\alpha,t} = l_\alpha + en_{\alpha,t} \tag{14}$$

where $n_{\alpha,t}$ indicates the number of agents using technology α in period t, and the parameter $e \in [0, 1]$ measures the strength of the externalities. Each time-step in the model consists of two stages that take place sequentially: the innovation stage and the decision stage.

Innovation stage

Technologies form a directed network where they represent the nodes, while the links express the generational relation. This network evolves in time due to the possibility that agents innovate and create new technologies. In every time-step, each individual agent innovates with independent probability p. Innovations take place differently depending on whether recombinations of existing technologies are allowed or not, something which is determined exogenously in the model.

- If recombinations are not allowed, each group of innovators using the same technology α branch out from α to jointly create a new technology α^* with intrinsic quality level $l_{\alpha^*} = l_\alpha + 1$. Thus, in this case the number of technologies created in time-step t equals the number of different technologies being used by the innovators.

- If recombinations are allowed, all innovators join up to create one single new technology β by recombining all the different technologies α_i they are currently using. The intrinsic quality level of the new technology is one unit higher than the highest quality of the technologies being recombined, i.e. $l_\beta = \max_i\{l_{\alpha_i}\} + 1$.

In either case, whenever a new technology is created, a link is formed from each of the technologies used by its creators to the newly created technology. If all the creators of a new technology come from the same technology we say that the innovation occurred by branching. Otherwise we say that the innovation occurred by recombination. Thus, it is clear that a technology has been created by recombination if and only if it has more than one incoming link in the network.

Decision stage

The decision procedure is different depending on whether the agent innovated or not. Naturally, all time-step innovators decide to use the technology they just created. Then, the agents who have not innovated in the current time-step synchronously decide which technology to use. The non-innovating agents will switch to the technology that provides them with the highest utility, once switching costs are taken into account. We assume that the cost of switching from technology α to technology β equals the geodesic distance $d_{\alpha\beta}$ between technologies α and β in the technological network. Thus, the benefit of switching from technology α to technology β is:

$$\Delta u_{\alpha\to\beta} = u_{\beta,t} - u_{\alpha,t} - d_{\alpha\beta} \tag{15}$$

A non-innovating agent using technology α will change technology if and only if there exists some other technology β such that $\Delta u_{\alpha\to\beta} > 0$. In that case the agent will choose the technology that provides him with the highest benefit $\Delta u_{\alpha\to\cdot}$. Ties are resolved randomly.

References

Arthur, B. 1989. Competing technologies, increasing returns, and lock-in by historical events. *Econom. J.* **99** 116–131.

Bruckner, E., W. Ebeling, M.A. Jimenez-Montano, A. Scharnhorst. 1996. Nonlinear effects of substitution. An evolutionary approach. *Journal of Evolutionary Economics* **6**(1) 1–30.

Carayol, N., J.M. Dalle. 2007. Sequential problem choice and the reward system in Open Science. *Structural Change and Economic Dynamics* **18**(2) 167–191.

David, P.A. 1985. The economics of QWERTY. *American Economic Review (Papers and Proceedings)* **75** 332–337.

van den Bergh, J. C. J. M. 2008. Optimal diversity: increasing returns versus recombinant innovation. *J. Econom. Behavior Organ.* **68** 565–580.

van den Bergh, J. C. J. M., P. Zeppini-Rossi. 2008. Optimal diversity in investments with recombinant innovation. Tech. Rep. 2008 -091/1, Tinbergen Institute.

Vega-Redondo, F. 1994. Technological change and path dependence: a co-evolutionary model on a directed graph. *Journal of Evolutionary Economics* **4**(1) 59–80.

Watts, D. J., S. H. Strogatz. 1998. Collective dynamics of 'small-world' networks. *Nature* **393** 440–442.