Sustainability and Industrial Change: The Hindering Role of Complexity[*]

Tommaso Ciarli[†] Karolina Safarzynska[‡]

November 29, 2020

To appear in:
Uwe Cantner, Marco Guerzoni, Simone Vannuccini (eds), *Handbook of Research Methods and Applications in Industrial Dynamics and Evolutionary Economics.* Edward Elgar

Abstract

A transition to a low-carbon economy requires moving to the production of goods that are less energy- and material-intensive than current practices. This may prove difficult, as producer objectives may not align with reducing pollution, unless this is a consumer priority, or is imposed by regulations. It has been argued that changing lifestyles and consumer preferences can drive technological change towards sustainability. In this paper we use the model by Windrum et al. (2009b) to show that the interactions between the populations of consumers, producers and technologies, when product components are interdependent, generate complexity, as a result of which changing consumer preferences may be insufficient to achieve sustainability objectives. Complexity may influence negatively the rate and direction of innovations towards the production of greener goods, causing a vicious cycle. Firms tend to remain stuck in local optima of the existing technological landscape, if most consumers are satisfied with the non-green characteristics of goods. As a result, firms are less likely to explore innovation possibilities to improve environmental performance of their products, which in turn reduces consumer expectations with respect to the environmental quality of future goods. As pro-environment consumers also imitate the higher preferences for non-green characteristics, firms have even higher incentives to improve those characteristics in the current technological paradigm than to explore new greener paradigms. The toy model proposed in this paper can be applied to study diffusion of 'green' products in a number of industries and to study environmental policies that can reduce complexity. The paper also offers a selected review of micro and industry level models of sustainable transitions.

Keywords: Sustainable transition; industry-demand co-evolution; interactions; complexity

[*] We thank the handbook editors for their support and comments and an anonymous reviewer from the SPRU Working Paper series for their comments. The model is reproduced from Windrum et al. (2009b). Ciarli has benefited from funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 649186 – Project ISIGrowth.

[†] Corresponding author: t.ciarli@sussex.ac.uk. SPRU, University of Sussex, Jubilee Building, BN1 9SL Falmer, Brighton, UK
[‡] Faculty of Economic Sciences, University of Warsaw, PL; ksafarzynska@wne.uw.edu.pl

1. Introduction

Agent-based models (ABM) have been recently suggested as a promising way of modelling climate change impacts, because they can accommodate the multiple and complex ways in which the economy and the environment interact (Farmer et al., 2015; Stern, 2016; Balint et al., 2017; Ciarli and Savona, 2019). Modelling disaggregated interactions between heterogeneous agents, structural change and innovations goes beyond the simplification of the Integrated Assessment Models (IAMs) that dominate in the climate policy assessment. IAMs have been criticized for adopting unrealistic assumptions (Pindyck, 2013; Stern, 2013; Lamperti et al., 2018). They rely on the constructs of equilibrium and representative agents. In such models, firms employ energy in production, which is a source of carbon dioxide emissions that accumulates in the atmosphere, reducing the productivity of output. On the demand side, a rational representative consumer decides each period on the allocation of income between savings and consumption. Savings are invested in capital accumulation, which then drives economic growth. The entire complexity of human decision-making on the supply and demand side is captured in the IAM by a single parameter: a discount rate.

In this paper, we argue that these models cannot deal with a changing structure of the economy in the context of low-carbon transitions. They underplay the role of interactions between consumers and producers and technological complexity in the process. To guide the economy towards sustainability, we need to understand how interactions between consumers and producers shape the direction in which firms innovate, and how boundedly-rational consumers choose between different products that differ with respect to their environmental impacts. In this paper we thus offer a toy agent based model to analyze these processes.

Our approach is motivated by the fact that transitions to a low carbon economy require a change in lifestyles and preferences of consumers to stimulate wider adoption of goods, which production and/or use are less energy- and material-intensive. This may happen because of consumers' inherent pro-environmental preferences or because they find 'green' goods more attractive, i.e. better and cheaper. Either way, consumers must have some motivation to purchase such goods. Their consumption choice, in turn, will influence firm behavior. This may induce firms to try to reduce the environmental impact of their goods and services, for instance by improving energy efficiency of their production techniques or by adopting less-polluting energy technologies. However, due to several interactions between consumers, producers and the technology, the process leading to a transition is not as simple as nudging consumers to buy energy-efficient appliances.

Interactions between consumers, producers, and the components of a technology play an important role in the process. Such interactions may lead to increasing returns such as economies of scale or learning by doing that can lock in a system to a single technology (Arthur, 1989; Cowan and Gunby, 1996). Once a technology becomes dominant, its consecutive adoptions enhance its leading position. Once this happens, it is difficult to change the direction of technologies progress (Scoones et al., 2007; Johnstone and Stirling, 2015). This can be best illustrated with the lock-in to fossil fuel technologies, and difficulties in promoting diffusion of renewable energy. In this context, a transformation of the dominant technological regime requires not only investments in R&D, or subsidies for new technologies, but also facilitating access to knowledge, technological possibilities and opportunities, or interactions among companies, scientists and engineers (Grin et al., 2010; van den Bergh et al., 2011).

On the supply side, technological progress is not straightforward. Innovations are inherently uncertain. For simple technologies, which can be decomposed into parts, each working independently from the other, the problem of improving the technological performance is relatively simple: it suffices to improve each component (Simon, 2002; Kauffman and Levin, 1987; Frenken et al., 1999). However, many products are complex, they are built of many interdependent components or sub-technologies. In this case, the functioning of one component depends on performance of other components. As a result, changing any component on its own influences the way in which all other components work and contribute to the technological performance. This is in particular problematic in the context of sustainability, when firms attempt to reduce environmental impact of a single component. For integrated products, even when the environmental impact of one component is reduced, its interactions with the other components of the technology may increase the overall environmental impact of the product. For example, using a less polluting fuel for a car, may require changes in the car's engine and the exhaustion pipe; otherwise the fuel may be more polluting than its predecessor.

Addressing the environmental impact of a good made of several integrated component may be harder when each is produced by a different company, which controls only part of the production process (Langlois, 2002; Ciarli et al., 2008). Each company may not be fully aware of the changes implemented in other components, how they integrate in the final good, and what are the consequences in terms of environmental impact.

On the demand side, the impact of technological choice is also uncertain. The rebound effect (Steve Sorrell, 2007; Barker et al., 2009; Stapleton et al., 2015) provides an example to illustrate how difficult it is to assess an environmental impact of consumers' choice. The use of less polluting (and energy intensive) goods may induce consumers to increase their use, with an overall increase in emissions. There is also high uncertainty about externalities and unforeseen effects. For example, the car was initially welcomed as a 'clean' option with respect to polluting horses (Windrum et al., 2009a). This could have been the case when only few consumers could afford and use low-energy cars. But as we know, cars became a major source of pollution when they diffused massively and their engine became more powerful and energy intensive (unforeseen at the time of the transition from horse to car).

As a result of the interactions between complex technologies and consumers behavior, the direction of technological change comes with advantages and disadvantages for environmental sustainability that are difficult to assess a priori (Leach et al., 2007). Advantages and disadvantages of innovations are difficult to assess a priori because of the full range of behavioral and technological effects that they may have. This difficulty can affect the behavior of consumers, influencing their preferences and incentives. Different consumers may form expectations about future innovation pathways, which influence their preferences and may affect firms strategies.

Let us illustrate such uncertainty with an example. As simple as it may seem, it is not obvious to compare the environmental impact of washing dishes by hand or using a dishwasher. A Google search of 'dishwasher vs hand washing environmental' returns millions perspectives and answers. This is because such comparison requires a life cycle analysis of all components used in hand and machine washing, i.e studying where the metal comes from, how it is assembled, or how the detergent is produced, as well as an assessment of how consumers employ each component and how each of them impacts on the environment. On top of this, small behavioral changes can influence producer choice and changes in technical features (e.g. energy efficiency and cost), which in turn influence consumer behavior; and so on.

Such technological and behavioral complexities have been largely ignored in climate policy discussions. They can be studied with evolutionary-economic models that employ an agent-based modelling (ABM) approach. Following the ABM approach, macroeconomic outcomes emerge from interactions between large numbers of distinct agents in distinct networks (Tesfatsion, 2006). In ABM, agents are modelled as independent entities having their individual objectives, preferences, knowledge, who perceive and adapt to changes in the environment. They are often described by rules that can accommodate a variety of boundedly rational behaviors, but also include rational behavior and utility maximization. The interactions between agents and the feedbacks from aggregate emerging outcomes, are the sources of nonlinear dynamics and of further emergent phenomena. Evolutionary ABM have proved capable of explaining a number of stylized facts, which traditional economic approaches rule out as 'out-of-equilibrium' properties such as the cascades of bankruptcies of firms and banks or business cycles. Such models have been widely adopted in modelling industrial dynamics and technological change (Malerba and Orsenigo, 1997; Janssen and Jager, 2002; Oltra and Saint Jean, 2009; Windrum et al., 2009b,a; Safarzynska and van den Bergh, 2010), economic growth (Dosi et al. (2010); Cincotti et al. (2010); Ciarli et al. (2018)), or the cascades of bankruptcies in financial markets (Tedeschi et al., 2012; Thurner and Poledna, 2013).

Over the last two decades, evolutionary ABM have achieved an increasing attention in modeling different aspects of sustainability transitions. For instance, the co-evolutionary models discussed in Section 2, have offered important insights on how to unlock the market, where evolving consumers preferences affect the direction towards which firms innovate. More recently, authors have combined evolutionary models with energy markets and/or climate modules (e.g. Gerst et al., 2013; Wolf et al., 2013; Ponta et al., 2018; Lamperti et al., 2018) to study interactions of different sub-systems in the economy and how they can generate a systemic risk, or can amplify damages from climate change.

In this paper we focus on, and extend, a toy-model by Windrum et al. (2009a) that explains how the interactions between consumers, between firms, between consumers and firms, and between technological components may influence the environmental impact of consumption and related production (Section 3). Before presenting and discussing the model, Section 2 provides a brief overview of the evolutionary-economic literature relevant to deepen our knowledge about sustainability transition. Section 4 concludes and proposes extension to the toy model.

2. A Selected Literature Review

Evolutionary-economic models can provide important insights to modelling sustainability transitions (Ciarli and Savona, 2019; see Safarzynska et al. (2012) for a review of policy oriented evolutionary-economic models; and Balint et al. (2017), Lamperti et al (2019) and Hafner et al. (2020) for overviews of evolutionary ABM). In this section, we discuss how technological change, evolving preferences and consumer-producer interactions (co-evolution) are modelled in evolutionary-economic theories, and discuss their relevance to understand sustainability transitions.

Industry dynamics models explain economic and organizational change as a result of evolutionary forces acting on the population of firms: innovations introducing new varieties to the population and selection causing differential growth of firms. In such models, heterogeneous firms actively search technological landscapes for better solutions or to imitate frontier technologies (Nelson and Winter 1982). New technologies and products can emerge at any time. Most early industry models depict products (technologies) as defined over one or two dimensions such as quality and cost. However, transitions to sustainability generally involve changes in large technological systems or complex technologies embodying many technical components, where different sub-technologies co-evolve. This creates a challenge as changes in one sub-technology, for instance improving the technical characteristic of a single component, may negatively affect the functioning of other components, reducing the overall performance of the technology. Examples of non-modular technologies are numerous: cars, aircrafts, or computers combine different technological solutions in a single product. A particularly well known way to represent interdependencies between sub-technologies is to use the NK-model originally developed in the context of biological evolution (Kauffman and Johnsen, 1991; Kauffman, 1993). It has been shown that as the complexity of technologies increases, as a function of the interdependence between its components, it becomes more difficult to find an element to be improved (Kauffman, 1993; Auerswald et al., 2000). Optimizing the performance of non-modular technologies is inherently difficult because the 'fitness landscape' consists of many local optima. Building on the concept of fitness landscapes underlying the NK-model, Alkemade et al. (2009) study transitions pathways as a series of changes in sub-technologies leading to a transition from a current system to a new (locally optimal) system. The authors show that due to the path dependent and irreversible nature of innovation in complex technologies, an initial transition step along some preferred path may cut off paths that later may turn out to be more desirable. The authors apply the model to study the possible transition paths in the mobility sector, by comparing the relative performance of potential future car systems as a result of incremental innovations in car characteristics.

The important insights from this line of research is that maintaining diversity of technologies options is important to prevent lock-in to a single technology that initially looks promising, but overtime may turn sub-optimal.

Diversifying investments in technological options allows also for combining existing technologies and ideas, which is widely recognized as an important source of innovation (Tsur and Zemel, 2007; Weitzman, 1998). Here, experimenting with variations of existing technologies may contribute to knowledge creation. However, maintaining the diversity of options is generally expensive for a single firm, and at the same time the benefits from each innovation are uncertain (Safarzynska and van den Bergh, 2010, 2013).

Zeppini and van den Bergh (2011) focus on the trajectory of technologies as an outcome of firm innovation. They extend Arthur (1989) lock-in model introducing the possibility of innovating by recombining technologies from different trajectories. The two competing technologies are green and brown, which are substitutes. The authors show that the recombination of the technologies may offer hybrid technological pathways, with lower environmental impact than that of incumbent technologies.

Most evolutionary models of industrial dynamics reduce the consumer side to a static selection environment, while assuming that the processes of innovation, creation, and selection are independent. Theories of ‘technological push’ emphasizes the role of market forces in the process of change. They rely on the one-way causal determination from science to technology and production, largely ignoring the role of economic factors in the process of change (Dosi, 1982). In turn, theories of ‘demand-pull’ assume that the market is capable of signaling consumer needs through the relative movements in prices and quantities and consequently of pulling the innovative activities of producers in a particular direction of search. Both approaches are criticized for offering a partial explanation of market dynamics and technological change. Many successful innovations, which seem to be unrelated to user needs (e.g. innovation emanating from blue-sky research), stem from user-producer interactions (Mowery and Rosenberg, 1979).

A number of evolutionary models have been proposed to study technological change as a result of the co-evolution of technologies on the supply side and of consumer preferences on the demand side. In models of demand-supply co-evolution, the substitution of an incumbent by a new technology relies on the pace of technological change and evolving consumer preferences. For instance, Windrum and Birchenhall (1998) propose a formal model of demand-supply coevolution to examine determinants of technological succession. In their framework, firms offer products to satisfy clients in consumer classes, to which they are randomly assigned. In addition, firms engage in product innovation to attract new consumers. Consumers move between consumer classes depending on the relative attractiveness of products offered by incumbent firms. They imitate the consumption choices of their peers, if this can help them achieve a higher utility. Evolving preferences determine which firms are successful, and thus the direction of product innovations.

Building upon this line of research, Safarzynska and van den Bergh (2010) propose a co-evolutionary model that conceptualizes five different mechanisms of increasing returns, which can prevent diffusion of new products. On the demand side, imitation is an important mechanism that, by exploiting information already acquired by others, allows saving on costs of individual learning, experimentation, or searching. Following others' choices may be the source of additional advantages, such as the creation of a network of users. In addition, advertising contributes to informational increasing returns: the better the product is known, the more individuals are willing to buy it, and the higher is the probability of creating a network of users. On the supply side, the product quality improves as firms increase their competence through production and market experience, referred to as learning-by-doing, and learning-by-using. Intuitively, the more a particular product is adopted, the more resources (R&D budget) are available for the product development and its quality improvement. Typically, incumbent firms have more resources to invest in promotion, which improves their market advantage. Finally, economies of scale cause the average cost of production to fall with the number of units produced. The model has been used to study a number of policy instruments aimed at escaping lock-in to a single technology. The analysis reveals that the effectiveness of such policies depends on the structure of network interactions between boundedly rational consumer as well as the strength of different types of increasing returns. The authors suggest how knowledge about feedback loops can be used to design policies to prevent the dominance of a single technology.

The key insights from the discussed literature for sustainability transitions are: creating a network of users of environmental innovations is crucial to promote their diffusion; investments in diverse technological options can prevent lock-in to technologies based on fossil-fuels and lead to the emergence of recombinant innovation with possibly improved environmental impacts compared to incumbent technologies. However, as we will show below, the complexity that emerges from the interaction between consumers, producers and technological component, may require a stronger intervention in directing technological change.

Other policies which can help prevent dominance of a less sustainable alternative have been identified in the evolutionary-economic literature as: labelling of environmental products; informing consumers about environmental performance of the products, or sharing ownership of products. For instance, Bleda and Valente (2009) show that environmental labels reduce uncertainty and drive production towards more sustainable goods. Buenstorf and Cordes (2008) argue that because green technologies are not necessarily better than brown technologies in terms of performance, promoting environmental values is more relevant for the diffusion of green goods than learning about the environmental impact of one product. Pasimeni and Ciarli (2018) show that where adoption of an environmental product is too expensive for a single consumer, creating coalitions of consumers who co-own the product may reduce production and material use in the economy.

3. A Toy Model

In this section we describe a model, first published in Windrum et al. (2009b), which helps to conceptualize the relevance of the interactions between consumers, producers and technology components, for the emergence of less polluting products. In Section 3.8 we extend the model to capture the uncertainty rooted in the technological change towards more sustainable goods. We add the interaction between several components of a technology, which makes the exploration of technological landscape complex, reducing the relevance of the expectations on future technological trajectories for consumer choice. The uncertainty for both producers and consumers increases with the complexity of the technology, as firms discover information about the technology while exploring it. Such uncertainty may not allow to fully exploit the technology green potential, if firms randomly start on a search path that leads to local optima, where the global optimum is the most sustainable technology in a given technological paradigm. The more complex and newer is the technology, the higher the chance for a firm to follow a suboptimal research strategy and lock-in in local optima; and the higher the chance for consumer to lower their expectations about the green potential of the new technology.

We use this model as it captures several features that apply to the co-dynamics between consumers and producers that are crucial to understand how firms improve the environmental impact of their goods, and the process of their adoption. Innovation in this model is the outcome of a co-learning process between producers and consumers. The model is also quite flexible: it can be easily extended to capture more sophisticated firm and consumer behavior, to add more sectors, such as finance or energy, and to include a macroeconomic account.

The model features two types of interacting agents: firms and final consumers. Firms produce a good with a vector of product characteristics that define its use properties (Lancaster, 1966a), a price and an environmental impact (from consuming it). Firms target a given consumer class, endowed with given preferences. Firms can improve the feature of the goods that they produce through innovation, which may affect its cost (therefore price), quality (the vector of characteristics), or the environmental impact of consuming it. Environmental impact in the model is a property of the good, which depends on its 'environmental fitness', rather than a property of the production process (as more commonly analyzed in the literature). Because pollution depends on the goods purchased, consumers are concerned about the pollution externality of using a given good, rather than about the technology to produce it. This is like assuming that the environmental impact of a good depends both on its production and use. Firms may face a trade-off between increasing the quality of the good (or some of the characteristics that define it), reducing the price (negatively related to quality) and decrease its environmental impact, i.e. pollution (related to quality in a way unknown to firms and consumers). For example, firms may increase a car's speed, which will also tend to increase pollution, or may focus on electric cars, which also decrease autonomy and increase the price. We assume that incumbent firms have an advantage with respect to new entrants in an existing technological paradigm, due to learning and accumulating knowledge. But new firms may enter new technological paradigms, once they are discovered (more on this below).

Consumers are distributed across different classes. By pertaining to a consumer class, they differ, among other things, with respect to their preference for the price, quality (along a vector of the product characteristics that define its use), and the environmental pollution caused by using a good. Within a class, consumers are homogeneous. This introduces in the model the crucial difference between individual and collective benefits of individual choices. The actions of a small number of environmentalists through consumption may have a small impact on the stock of pollution, unless their action is imitated by similar consumers. Two opposite outcomes may occur: classes of environmentalist consumers manage to attract consumers that are initially less concerned about the polluting features of the consumed good. Or environmentalists are so poorly catered by existing companies, that they may need to change class, give up their environmental preferences and imitate consumers that enjoy higher utility by caring more about non-green product characteristics than about pollution.

Technological change is then the consequence of consumers and firms interactions. Consumers purchasing behavior signal their preferences, and influence firms innovation behavior. Changes in the produced goods, as an outcome of innovation, modifies purchasing behavior. This is because purchasing behavior and innovation influences the dynamics of both populations. The population of firms changes depending on how successfully they innovate, where success depends on the preference of the prevalence of consumers, and the ability to capture a sufficient large niche. Consumers population changes as a function of the product available on the market: consumers in a class that enjoy high utility (because preferences are better aligned with the prevailing goods) will thrive and outnumber consumers in classes whose preferences are poorly matched by firm innovation. The model thus features co-evolution between the two populations, influencing both their dynamics, and changes in the technology.

The product features (price, quality characteristics and environmental impact) are defined over a given paradigm. A paradigm shift implies a shift in the boundaries of these features: price, quality and environmental impact. That is, improvements of the quality and the environmental impact of a good within a paradigm are limited (incremental innovation); a paradigm shifts, instead, opens up new frontier for radical and incremental changes.

Within a technological paradigm, firms and consumers know the boundary of the product environmental impact, i.e. the minimum pollution that it can cause, if firms innovate and move towards that boundary. Because both firms and consumers know the attainable environmental impact in a given paradigm, consumers assess firms based on their relative achievement with respect to this known minimum environmental impact.

Pollution is a function of the vector of product characteristics that define its quality. For instance, a bicycle is slower, less comfortable in long distances, and more exposed to the weather than a car, but it has a lower environmental impact. When innovating, firms change both the characteristics of a product that cater specific user services (e.g. the speed of mobility) and the environmental impact. In general, we assume that increasing the use characteristics of goods, increases their environmental impact when consumed. A good's environmental impact is thus modelled as a technological landscape of several dimensions, where each dimension is one of the product characteristics. All characteristics contribute to a good's pollution: the impact of one characteristic on pollution is given by the position of each of the other characteristics on the technological landscape. This may make environmental

innovations complex, depending on the interaction between the different product characteristics.
The level of complexity determines the uncertainty of technological change aimed at reducing the environmental impact of the consumption of goods. This implies that each innovation may cause several unintended consequences. In other words, certain directions of technological change may result in more pollution than expected, even when the innovation is believed to reduce environmental impact.
Through successive innovations, economies may move to a new paradigm. For this to occur, firms need to find their way through the technological landscape towards its maximum, i.e. the minimum pollution that a product produced within that paradigm can cause. Once they achieve a paradigm's global optimum (minimum pollution), firms open up opportunities to explore a new paradigm, with a new maximum, i.e. lower minimum pollution. Moving to new paradigms, in our model, is then crucial to reduce the environmental impact of consumption. When firms enter a new paradigm, however, they face a radically new, unchartered and complex technological landscape, which they need to explore in order to improve the product characteristics and environmental fitness.
The model was developed based a number of empirical regularities (Windrum et al., 2009b). On the demand side: wealthy consumers seek for different options when negative externalities are too high (such as pollution in large cities or the congestion of public transport); consumers derive their utility from a set of product characteristics, beyond price, that define the use of the good and motivates its consumption (Lancaster, 1996); consumers have limited information and different beliefs about future technological feasibility (Archibugi, 2017); consumption choices generate externalities for other consumers. On the supply side: new technological paradigms retain features of older paradigms ("deep path dependency"); new firms champion new technologies (Klepper, 1996); pollution externalities induce environmental innovations (Safarzynska and van den Bergh, 2010); firms target different niches of consumers.
In the reminder of this section we describe the toy model in detail (3.1-3.6), summarize its main properties (Sec 3.7), and then extend it to analyze how technological complexity and uncertainty may influence pollution through firm and consumer behavior (Sec 3.8). The model is in discrete time. We suppress the use of the time index for clarity, unless when it is needed to distinguish between current and lagged variables.

3.1. Demand

We model $\hat{C} = 500$ consumers distributed across $N = 20$ classes j. Consumers are heterogeneous across classes, but homogeneous within classes: preferences differ across classes, and when a consumer changes class, they also change their preferences and budget constraint. In $t = 0$ consumers are distributed equally across classes (in each class there are 25 consumers).
A class utility (u_j), if function of the good's price (p_i), a vector of use characteristics $(\vec{x}_i)$, and the environmental impact generated by consuming it $\left(s(\vec{x}_i, G)\right)$, which is in turn a function of

the characteristics ($\tilde{x}_i$) and of the stock of pollution G. Each firm produces an heterogeneous good (Section 3.2), therefore we index a good's feature with that of the producing firm i. Formally, a class utility is expressed as:

$$u_{l,j} = v(m_j, p_i) + d(\vec{x}_i) + e(s(\vec{x}_i)) \quad (1)$$

where m_j is the budget constraint of all individuals in class *j*. The three terms of the class utility function have the following form:

$$\begin{aligned} v_j &= \alpha_j \sqrt{m_j - p_{t-1,i}} \ \ \forall p_{t-1,i} < m_j \\ d_j &= \textstyle\sum_{h \in z^j} \beta_{j,h} \sqrt{x_{t-1,h,i}} \\ e_j &= n_j \frac{\left[E_j(s_{t-1,i}) - \underline{s}\right]^{1-\rho}}{1-\rho} \forall \underline{s} < E(s) \end{aligned} \quad (2)$$

where α_j and $\beta_{j,h}$ are the consumer preferences with respect to the price and quality of the good (determined by a vector of characteristics $\vec{x}_i$).

The first component of u_j simply represent a consumer preference for saving (in a given class *j*). The price of the good p_i is relatively more relevant the lower is the consumer budget constraint. In other words, the preference for savings decreases with the budget constraint: consumers in wealthy classes are less influenced by prices in their purchasing decision.

The second component is the direct utility from consuming a good and benefiting from its h^{th} characteristics (as in Lancaster, 1966a; Saviotti and Metcalfe, 1984; Gallouj and Weinstein, 1997). For example, speed, memory, hard drive, and weight of a laptop. Note that each service enters equally in the consumer utility, and the contribution of each characteristic ($\beta_{j,h}$) may differ across classes. In other words, the direct utility is a weighted average of the level of a good's characteristics, where the weights are the preference terms $\beta_{j,h}$. Characteristics change across technological paradigms (Sections 3.4 and 3.5), to reflect the fact that a laptop is radically different from a typewriter, or a digital technology is radically different from an analogue one. All products produced within a paradigm *z* have H_z use characteristics, which yield direct utility to consumers when the product is consumed (Lancaster, 1966a,b; Saviotti and Metcalfe, 1984).

The third component, the environmental utility, is a composite function that reflects the hyperbolic absolute risk aversion (HARA) of consumers (e.g. Merton, 1971) towards the negative externalities of environmental pollution *G*, where ρ and η_j are parameters that reflect, respectively, the relative risk aversion toward pollution and the discount rate of a consumer class; $E_j(s_i)$ is the consumer class expectation of firm environmental fitness, which is related to the knowledge that consumers have about the technology (and thus includes uncertainty); and ($\underline{s}$) is the minimum level of environmental fitness that a class is ready to accept from a good, and the firm producing it. Note that ($\underline{s}$) differs from the level of pollution, which depends on what other consumers purchase, not only on the single good. In other words, utility increases as firms produce more environmentally sustainable good, for a given discount rate and risk aversion. It should be noted that when we refer to a firm environmental

performance, in this model we only refer to the features of the good they produce, and not to a firm production process.
The expected environmental fitness of a firm (i.e. of the product produced) $E_j(s_i)$ is a combination of the fitness of the best technology available in the market ($\hat{s}(z)$) in a given time period (*t*)[1] and the firm environmental fitness $s(\vec{x}_i)$:

$$E_j(s_i) = \eta_j^p \frac{s(\vec{x}_i)}{1+\hat{s}(z)-s(\vec{x}_i)} \tag{3}$$

where $\eta_j^p \in [0,1]$ is a weight that consumers attach to the current environmental impact of design i[2] relative to the technological promise of the most recent paradigm $\hat{s}(z)$ (note that a design ($\vec{x}_i$) is specific to a firm *i*).
Finally, the minimum level of environmental fitness (higher fitness means lower impact) that a class is ready to accept in a good ($\underline{s}$) is a logistic function of the pollution stock:

$$\underline{s}_t = \begin{cases} \underline{s}_\tau + \dfrac{\hat{s}_t/2 - \hat{s}_\tau}{1 + \left[\dfrac{\hat{s}_t/2 - \hat{s}_\tau}{s_0} - 1\right] e^{-r(G_{t-1}-G_\tau)}} & \Delta G_{t-1} > 0 \\ \dfrac{\underline{s}_\tau}{1 + \left[\dfrac{\underline{s}_\tau}{s_0} - 1\right] e^{-rG_{t-1}}} & \Delta G_{t-1} < 0 \end{cases} \tag{4}$$

where τ is the time period in which the boundaries of the minimum environmental fitness may change due either to a change in the paradigm *z* – which implies a change in $\hat{s}_t$ – or to a change in the sign of pollution growth (ΔG); *r* is the rate of growth of the minimum level of environmental fitness of a good with respect to pollution; and s_0 is the lower asymptote. As pollution increases, we assume that consumers may become more demanding with respect to the environmental fitness of the good that they would like to purchase. Consumers are assumed to be relatively less concerned about growth in pollution when pollution stock is low, than when it is high: this feature is captured by the logistic shape.
This specification captures three key properties about consumer preferences with respect to environmental sustainability. First, the role of expectations: in the initial phases of diffusion, consumers need to make their judgement based on what is known about the environmental impact of a good (e.g. from research), rather than from the actual observed impact on the environment caused by the use of goods. Second, the expectations about how much a technology may pollute, does not always match its actual impact, once the good is diffused (as this depends on the dynamics of the consumer populations). Third, different consumers, with different risk aversion and care for future generations, weight the relevance of expectations and observed pollution differently.
Consumers may change preferences due to 'imitation'. If they observe that other consumers, in different classes, are enjoying a relatively higher utility, they may change class. This consumer population dynamic is oversimplified in this model, and can be easily extended to better capture social imitation. In its current form, we assume that classes in which consumers enjoy relatively higher utility attract consumers from classes where consumers experience a

relatively lower utility. In other words, a class that is well catered by existing goods (i.e. goods that balance the trade-offs between the direct, indirect, and environmental preferences of that consumer class), experiences a higher average utility than a class that is not well catered for by the existing goods.

Formally, the movement of individual consumers across classes is modelled as a replication dynamics. Classes with above-average utility, grow as a proportion of the total population, while classes with below-average utility decline. As a result, the combination of preferences in the population also change, moving towards the preference of the classes that grow in number of consumers (the total population is fixed). In turn, this change in consumer population (and average preferences) also changes the signal for firms, which may need to adapt their innovation behavior to accommodate the changing distribution of consumer preferences. Because with a pure replicator dynamics only one class is likely to survive in the limit, which would also lead to a single dominant design, and a single firm dominating the whole market, we use a 'tamed' replicator (Wirkierman et al. (2018)): an intensity parameter f tempers the strength of selection, allowing a number of classes with similar utility to have the same share $\psi_{j,t}$ of total consumers $\hat{C}$.

The number of consumers $C_{j,t} = \psi_{j,t}\hat{C}$ in each class *j* is computed as a ratio $\psi_{j,t}$ of the total number of individual consumers:

$$\psi_{j,t} = \psi_{j,t-1}\frac{\overline{u}^{f}_{j,t-1}}{\overline{U}^{f}_{j,t-1}} \tag{5}$$

where $\overline{u}^{f}_{j}$ is the average utility of class *j*:

$$\overline{u}^{f}_{j,t} = f\,\frac{\sum_l u_{l,j,t}+\varepsilon^u/C_{j,t}}{\sum_j \psi_{j,t}(\sum_l u_{l,j,t}\varepsilon^u/C_{j,t})} = round\left[f\,\frac{\overline{u}_{j,t}}{\overline{U}_t}\right] \tag{6}$$

$\overline{U}^{f}_{t-1}$ is the average utility across all classes; $u_{l,j,t}$ is the utility of a single consumer *l* in class *j*; and ε^u is a small parameter allowing each class to survive through time, so that it can be populated again, in case it becomes attractive when its fitness change (e.g. because of a change in the technological paradigm).

In each time period, consumer classes access the market in random order (a different one in each period). When it is their turn, each consumer in a class select the firm that best satisfies their utility. To simplify, we assume that each consumer buys one unit of the selected good. Firms use their inventories and finished goods to match the demand from a class. When they run out of inventories, consumers move to the second best firm, and so on, until all consumers from the class have purchased one unit.[2] As firms run out of inventories, it is possible that a consumer class finds no firms that can offer a good that attain a utility that is larger than the utility from not consuming $\left(u_j < \alpha_j\sqrt{m_j}\right)$. Similarly, consumers from a given class my not find a firm that sells at a price which is below their budget constraint. Finally, firms may produce below the overall demand, leaving consumers at the end of the line in a given period with no purchasing options. When consumers do not consume for one of these reasons, their utility comes from saving, or consuming the budget on a different market: $\alpha_j\sqrt{m_j}$.

3.2. Supply

We model *F* firms indexed by *i* producing an heterogeneous good, with different use characteristics, to satisfy one unique consumer need. Firms are initially homogeneous, endowed with the same market share and capital, the only factor of production. Production is kept to its simplest form, to allow focusing on the innovation process, industrial dynamics, and the interaction with consumers. As times goes by, firm market shares depend on the relation between consumer preferences, the price, quality and environmental fitness of the produced good. To produce the good firms invest in capital, which defines their production capacity. Depending on the relation between production and demand, firms accumulate non-perishable inventories, which are carried on from one period do the next. Firms innovate in order to improve their good, but depending on the market signal they receive from the consumers buying from them, they may follow different innovation paths in the technological landscape. Firms that do not manage to maintain a sufficient amount of capital, exit the market.

Firms define a target level of output (y_i^*) as a linear combination between consumer demand ($D_{t,i}$) and actual sales ($S_{t,i} = min(D_{t,i}, q_{t-1,i})$), which cannot be higher than the available inventories $q_{i,t-1}$:

$$y_i^* = \lambda^y D_i + (1 - \lambda^y) S_i \tag{7}$$

where $\lambda^y \in [0,1]$ allows to adjust smoothly to changes in demand and avoid sudden oscillations.

Given y_i^* and the financial constraint, a firm may (dis-)invest, according to the following rule:

$$I_{t,i} = \begin{cases} \lambda^c \min\left(y_{t,i}^* - k_{t-1,i}, w_{t,i}^*\right) & \text{if } y_{t,i}^* > k_{t-1,i} \\ -\lambda^c \min\left(k_{t-1,i} - y_{t,i}^*, k_{t-1,i}\right) & \text{if } y_{t,i}^* < k_{t-1,i} \end{cases} \tag{8}$$

where $\lambda^c \in [0,1]$ represents potential physical constraints in changing production levels in the short run. Firms invest when the target output is above the available capital, otherwise they disinvest and sell capital. When they invest, the amount is the minimum between the capital needed to achieve the desired level of output, and the financial constraint, the sum of the cumulated financial resources and last period profits: $\left(w_{t,i}^* = w_{t-1,i} + \pi_{t,i}\right)$. When output decrease and a firm needs to disinvest, they sell the difference between the available capital and the capital required to produce y_i^*– unless the difference is larger than the available capital, in which case they sell only the remaining capital available.

Changes in the capital stock then depend on the above investment rule and the financial resources in *t*:

$$k_{t,i} = \begin{cases} k_{t-1,i} + I_{t,i} & \text{if } w_{t,i}^* > 0 \\ \max\left(k_{t-1,i} + w_{t,i}^*, 0\right) & \text{if } w_{t,i}^* < 0 \end{cases}$$

(9)

As a result of the investment, the stock of financial resources that will be available in the following periods also changes:

$$w_{t,i} = w^*_{t,i} - \left(k_{t,i} - k_{t-1,i}\right) \tag{10}$$

Profits also form part of the financial assets available to firms to invest in the following period. They are computed as the difference between monetary sales (sales times unit price of the good) and costs (output times variable and fixed costs):

$$\pi_i \ = \ S_i p_i \ - \ c_i y_i \tag{11}$$

Price is given by a simple mark-up rule on costs, where the mark-up is fixed (equal across firms): $p_i = (1\ +\ \nu)c_i$. Costs depend on a fixed component *F* and a variable component (c_h) that depends on the quality of the good, i.e. the characteristics *h* that define it

$$c_{t,i} = \frac{F}{1+y_{t-1,i}} \sum_{h \in Z_i} c_h x^2_{t,i,h} \tag{12}$$

where z^i is the set of use characteristics that define the good. We assume that each characteristic *h* has a given cost c_h: improving the quality of *h* also makes the good more expensive, more than linearly. As we discuss in Section 3.4, firms face a trade-off between quality and price when they innovate. $y_{t-1,i}$ is the level of production in the previous period: we assume that there are some positive returns to scale, due to learning, which contribute to reduce the cost as firms gain market share, decreasing the trade-off between price and quality. Finally, firms produce using a production function with constant returns and capital as the sole input

$$y_{t,i} = \ k_{t-1,i} \tag{13}$$

Output is used to cumulate inventories to be sold in the next period

$$q_{t,i} = \ y_{t,i} - \ S_{t,i} + \ q_{t-1,i} \tag{14}$$

3.3. Environmental impact of goods

The environmental impact of using a good depends on the environmental fitness that the firm producing it reaches. Goods with high environmental fitness are less polluting than those with low environmental fitness, so the more consumers purchase goods with high environmental fitness, the lower is the rate at which the stock of pollution grows. We assume that the environmental fitness depends on the quality of each of the use characteristics of a good. For example, the environmental impact of a car may depend on it speed, size, and acceleration. We thus compute environmental *fitness* as the average fitness of the use characteristics $\vec{x}_i$ over all characteristics H^i_z that form part of the good in paradigm z^i:

$$s(\vec{x}_i) = \frac{\sum_{h \in z^i} \varphi_{z,i,h}}{H_{z^i}} \tag{15}$$

where $\varphi_{z,i,h} H_{z^i}$ is the fitness of the single characteristic; and H^i_z is the set of characteristics that defines a good in paradigm *z*. In this basic version of the model we assume that the technology is modular, i.e. each characteristic contributes to the environmental fitness independently. When this is the case, firms can improve each characteristic independently from the others. The choice to innovate in one or the other direction, is driven by the trade-off between improving the characteristic and improving its environmental fitness (as we discuss in Section 3.4 below).

The environmental impact of consuming the good produced by firm *i* is a decreasing function of environmental fitness, with a steeper slope for intermediate levels of fitness:[3]

$$\zeta_i = \frac{\hat{\zeta}}{1+\left[\frac{s(\vec{x}_i)-s_0}{\phi}\right]^2} \tag{16}$$

where $\hat{\zeta}$ is the maximum environmental impact of a good; s_0 is the minimum level of fitness attainable; and φ is a parameter that defines the rate at which an improvement in the environmental fitness of the good reduces its impact on pollution. The function is similar to a logistic. In the beginning, innovation is exploratory and yields marginal improvements to the environmental fitness: for very low levels of fitness, a fitness increase has a small impact in reducing the pollution impact of the good. As R&D activities continue, innovation manages to make larger steps, and improvement in fitness reduce the impact of using the good on the environmental sustainability. As the fitness reaches closer to its maximum, i.e. its maturity, returns to R&D to reduce the impact on the environment slow down. In other words, although increases in fitness are perceived by consumers in the same positive way, their impact on pollution differs for different phases of the innovation process, which come with different opportunities.

3.4. Innovation

As explained in Section 3.1, consumers choose goods depending on their utility, which depends on three features of the good produced by firms: the price, the vector of characteristics, and the environmental impact caused by its use. Firms have an incentive to reduce the price, increase the quality of its characteristics, and increase the environmental fitness. But they face trade-offs.

We assume that all firms undertake R&D in each period to modify the characteristics of the produced good, within a given paradigm. Modifying a characteristic has three effects: (i) changes the quality of the good, (ii) its cost (see Eq. 12) and (iii) the environmental fitness (see Eq. 15). We model innovation in two steps. In the first step, firms invest in R&D to innovate ('mutation'), attempting to change one characteristic. In the second step, firms assess this change, taking into account the preferences of the consumers in the class that they target (assigned at the outset and fixed throughout the firm's life time), and how the change modifies the trade-offs between quality, cost and environmental fitness ('evaluation'). Firms decide whether to retain the innovation(s) depending on the expected changes in the demand of the consumer class that they target.

Mutation

For simplicity we assume that R&D does not depend on firm revenues. All firms attempt an innovation, in each time period, on one random product characteristic *h*. There is a small probability ι that the innovation is successful and results in a mutation of the position $x_{i,h}$ of characteristic *h* on the technological landscape. When the innovation is successful, the firm draws a random number from a Standard distribution that defines the extent of the change of characteristic *h*:

$$\Delta x_{i,h} = N(0,1) \cdot \xi \tag{17}$$

where ξ is a parameter that allows to measure how local is the innovation process. If successful, as a result of R&D a one bit mutation then occurs: a change in the value of $x_{i,h}$ by a factor $\Delta x_{i,h}$.

Evaluation

If R&D was successful for at least one characteristic, the firm evaluates the environmental fitness of the new product (see Eq. 15), and simulates the impact of the innovation on the utility of its target consumer class *j* through the cost, quality and environmental fitness. For simplicity, we assume that firms has perfect information about the utility function of its target consumer class (see Eq 1-4). This is like assuming that firms invests in market surveys to elicit its consumers acceptance of an incremental innovation.

The final value of the product characteristic, following a successful innovation, then depends on the result of the evaluation process. If the utility of the targeted class increases as a result of the innovation, the final value integrates the change obtained with the mutation. Otherwise, the characteristic remains unchanged. Formally:

$$x_{t,i,h} = \begin{cases} x_{t-1,i,h} & \text{if } \overline{u}_j^{\Delta x} > \overline{u}_j \\ \min\left(x_{t-1,i,h} + \Delta x_{i,h}, 0\right) & \text{if } \overline{u}_j^{\Delta x} \leq \overline{u}_j \end{cases} \forall h \tag{18}$$

where $\overline{u}_j^{\Delta x}$ is the utility that consumers in class *j* would attain, should they buy in period $t+1$ from firm *i* the good with the modified characteristic value.

3.5. Paradigm shift and Pollution Stock

Goods pertain to a given technological paradigm *z*, which defines the boundaries of its use characteristics, their relation to environmental fitness, and the maximum environmental fitness that can be reached. To decrease the impact of consumption on pollution, firms need to move to a new paradigm, which has a higher potential fitness.

Because consumers evaluate the environmental fitness of a firm with respect to the environmental expectation that can be reached within a paradigm, by moving to a new paradigm firms have renewed opportunities to gain market share by increasing the fitness, but initially they may lose market shares, as they will be seen less performing with respect to firms still producing in older paradigms that have a lower maximum fitness (and expectations).

Incumbent firms are also constrained by their target class: we assume that a firm cannot switch to a consumer class when moving to an alternative paradigm. New firms are then crucial in opening new technological opportunities by exploring new technological paradigms, when they become available.

As discussed in the previous section, when they innovate, firms attempt to change the use characteristics of a good in order to improve its quality and environmental fitness. Their target is the maximum environmental fitness of a paradigm: $\hat{s}(z)$. They attempt to move in this direction on the technological landscape. This happens when product innovations towards higher environmental fitness improve the utility of a consumer class – given the trade-off between environmental fitness, cost and use value. When at least one firm reaches the maximum environmental fitness, we assume that this opens a window of opportunity to start researching for a new, improved, paradigm. We assume such a search to be exogenous (for instance driven by public sector research (Wirkierman et al., 2018)).

More formally, the exogenous search for a new paradigms starts when $\hat{s}(z) - s(\vec{x}_i) < N(0, \sigma^Z)$ for any firm *i* in time *t*. That is, when a firm reaches, for all product characteristics, an area of the landscape that is within a distance defined by a Normal distribution with average 0 and variance σ^Z.[4] When this occurs, exogenous scientific research identifies a new paradigm *z* after $\tau^z \sim U[\tau^z_{Min}, \tau^z_{Max}]$ time steps.

The new paradigm is defined by a higher maximum, potential, environmental fitness:

$$\hat{s}(z)_{t,z} = \hat{s}(z)_{t-\tau_z,z} + \Delta^{\varphi} \tag{19}$$

where Δ^{φ} is a parameter that measures the exogenous technological progress with respect to environmental sustainability.

Although the increase in the potential fitness can be exploited by firms to increase the utility of their class, the new paradigm is defined by a new, unknown, technological landscape, with a different relation between the product characteristics and environmental fitness, which firms need to explore:

$$\hat{\chi}_{t,z} = N\left(\hat{\chi}_{t-\tau^z,z}, \sigma^{\chi}\right) \tag{20}$$

where σ^{χ} is a measure of the technological distance between the old and the new paradigm. We assume that the number of the product characteristics x_h between paradigms is constant (H_z) and that two consecutive paradigms share at least one characteristic: $z_t \cap z_{t-\tau^z} \neq 0$. Therefore, the number of new characteristics that can emerge when a new paradigm is achieved is $U(1, H_z - 1)$.

We also assume that only start-ups can explore the new paradigm by way of their R&D activity. Due to lock-in, incumbent firms compete on the older paradigm.

Finally, the stock of pollution cumulated through periods depends on the fitness of firms products and on the number of consumers purchasing them (consumer preferences):

$$G_t = G_{t-1} + \sum_i \zeta_{t,i} \cdot S_{t,i} \tag{21}$$

A rate of pollution decay is implicit in the relationship between environmental fitness and impact (eq 16).

3.6. Market dynamics

Every $\tau^r \sim U[\tau^r_{Min}, \tau^r_{Max}]$ the least efficient firms (that do not produce anymore) and empty consumer classes (due to migration of consumers from one class to another (see eq 5)) are replaced as described below.

Consumers dynamics

To maintain a variety of consumer preferences in τ^r, any class *j* that is populated by less than $\delta^c = 2$ consumers $C_{\tau^r-1,j}$ is replaced by a new class j^{τ^r} with the same number of consumers but different preferences for the product characteristics: $\beta_{\tau^r,j,h} \sim U[\beta_{Min}, \beta_{Max}]$. When a class is replaced after the introduction of a new paradigm, we assume that its consumers positively value the product characteristics of the new paradigm $(z_{\tau^r} > z_{t-\tau^r})$. As a result, through time, a number of new classes that replace classes that were not well catered by existing goods, establish a market for the good produced by start-ups in the new paradigm.

Firms dynamics

In τ^r all firms whose capital stock $k_{\tau^r-1,i}$ is below a given level $\delta^k = 0.2$ are replaced by new entrants. New firms start with a capital stock and financial wealth equal to the market averages in $t-1$: $\left(k_{\tau^r-1,i} = \overline{k}_{\tau^r-1,i}\right)$ and $(w_{\tau^r,i} = \overline{w}_{\tau^r-1})$. They are also endowed with an inventory of finished products to satisfy consumers demand: $(q_{\tau^r,i} = k_{\tau^r,i})$. New firms also enter with improved product characteristics $x_{i,h}$ with respect to the incumbent, as if they were successful innovators: $\Delta x_{i,h} = N(0,1) \cdot \xi$.
Following the discovery of a new paradigm, all new firms adopt the new technological paradigm, if there is at least one consumer class that has entered after the new paradigm has emerged, and that they can target. The old paradigm's product characteristics $x_{i,h,z_{t-\tau^r}}$ are replaced by characteristics pertaining to the new paradigm: $x_{i,h,z_{t-\tau^r}} \sim U\left[\underline{x}_h, \hat{x}_h\right]$; where $\underline{x}_h$ and $\hat{x}_h$ are respectively the minimum and the maximum value of the product characteristics currently in the market.[5] Finally, the new firms randomly target a new consumer class j^{τ^r} purchasing products of the new paradigm.

3.7. Properties and main results

We summarize here the main general properties and results of the model. Details can be found in Windrum et al. (2009b,a). Details about the initialization of the model are provided in Appendix A.
In the basic version of the model the pollution stock (e.g. GHG emissions) grows at a decreasing rate. This is due to consumers tolerance with respect to pollution, which drives firms to increase the environmental fitness of their good – despite a relative reduction in the direct utility, and explore new technological paradigms. A pollution threshold in consumer tolerance is fundamental to push firms toward new, less polluting, paradigms: if consumers become

alerted to pollution, when close to their tolerance threshold, marginal changes in fitness have a large impact on utility, offsetting reductions in the product characteristics. Consumers enjoy a higher utility, on average, with less performing, more environmental goods.

In the basic version of the model we also observe market concentration, with the economy converging to oligopoly, even when firms compete on several co-existing paradigms. This is partly driven by consumers concentrating in few classes, reducing market differentiation. However, demand concentration is not a necessary condition in our model, as firms manage to target different classes with the same technology.

As expected, increasing the average relevance of environmental preferences across consumer classes reduces pollution. However, in our model this also has a perverse effect. Because the environmental component of the utility function is conditional on the potential environmental fitness of a technology (in a paradigm), for extremely high average environmental preferences firms may be better off exploiting the current paradigm and increase the value of product characteristics, rather than moving to new paradigms, where they will be punished for being too far from the potential frontier. In other words, if consumers expect a high environmental performance from a new technological paradigm, and they also have high preferences for environmental fitness, no firm has an incentive to move to the new technological paradigm, because by the time they manage to introduce incremental innovations, they would not be able to compete with firms performing at the edge of the older paradigm. This sounds familiar with many experimental green technologies, that require public support to attract private investors.

Aside from the average preferences, for a given low level of average environmental preferences across classes, higher heterogeneity of preferences across classes also reduces the pollution stock. This is because consumer classes with high environmental preferences, on average, attract more consumers, as they enjoy a higher utility when firms increase the environmental performance of their good. 'Eco-warriors' experience a larger utility, attract consumers that are less sensitive to the environment, increasing the demand for more eco innovations. When compared, the average environmental preference has a stronger impact on reducing pollution in our model, than the heterogeneity among consumer classes.

The model also shows that the positive effect of environmental preferences occurs when consumer preferences for product characteristics are sufficiently low. When there is a high trade-off between the use characteristics and the environmental fitness of a good, the former may prevail, reducing firm incentive to innovate towards environmental fitness. The preference for the product characteristics play a negative role on pollution abatement also when the average is relatively low, but the heterogeneity across consumer classes is high. With a very heterogenous population (with respect to their preferences for the use characteristics), firms have the option to focus on either the use characteristics or the environmental fitness of their good, which holds back environmental innovations.

With respect to the willingness to pay for improved environmental fitness, we find that the level of pollution depends on the distribution of consumer preferences with regards to the trade-off between environmental and price preferences. The larger the difference between price and environmental preferences, the higher the level of pollution. In other words, in the presence of consumer classes that are highly sensitive to price differences (high price elasticity), even the presence of consumer classes highly sensitive to pollution does not help reducing the environmental impact of consumption. When this is the case, firms target two different niches of consumers with old (low price and more polluting) and new technologies (high price and less polluting). The presence of the class of environmentally sensitive consumers, with their quota of green consumption, help maintaining pollution to a level that is low enough to allow firms to keep producing polluting goods for classes that prefer (or can afford only) cheaper goods.

3.8. Model extension: coordination and technological complexity

So far, we have assumed that the environmental impact of each product characteristic is perfectly modular. That is, it suffice for a firm to increase the environmental fitness of one characteristic to increase the overall fitness of the product. This is like assuming that improving the environmental fitness of goods is an easy walk on the technological landscape for all firms, and that there are no trade-offs between increasing one or the other characteristic, nor unintended consequences. Results above are driven by firm behavior as a pure response to market incentives under no constraints related to technological complexity.

In reality, improving the environmental fitness of a product characteristic, may come at the cost of other characteristics becoming more polluting. Biofuels may be overall less polluting than oil, but they also induce soil depletion. When they were first introduced, cars polluted cities less than horse manure, while they were slow and few. A dishwasher may be less polluting than hand washing, or a different dishwasher depending on several characteristics such as the relative use of water, material used and their provenance, organization of production (e.g. transportation involved), and so on. Batteries may allow to reduce the waste in energy production, but they impose a heavy demand on lithium extraction, refinement and transportation. In other words, while firms and consumers may have expectations on the potential achievement of a technology, they have only partial knowledge about how to reach it because of their complexity. It is also possible that intermediate improvements in the technology are less sustainable than original technologies, because of the partial improvement of some characteristics, that make the other characteristics more polluting.

The uncertainty related to the complexity of the potential impact of goods on pollution is magnified by the segmentation of production across different producers, which provide different product components, which are assembled by a final good producer. Producers of consumables face a coordination problem: they must coordinate producers of a vertically integrated industry, and usually have competences on specific components. The less modular is a good/technology, the more final producers may struggle to coordinate specific features of the final good, such as environmental fitness (Ciarli et al., 2008).

In sum, the lower is the product modularity of a technology/good (with respect to pollution), and the higher the division of labor to produce it, the more difficult it is for one single firm (e.g. the final producer downstream) to find ways to improve the environmental fitness of the produced good, and meet the expectations that consumers may have within a technological paradigm.

On the demand side, as explained in Section 3.1, consumers value the quality of a product characteristics, price, and pollution; they imitate better off peers and their preferences; and they have expectations on the potential environmental fitness of a good in a paradigm.

Depending on the initial distribution of preferences (as summarized in Section 3.7), consumers may steer firms towards different directions of innovation, more or less green. When we introduce technological complexity, though, the coordination between consumer preferences and firms decision becomes less obvious. Uncertainties about the impact of improving product characteristics on its polluting performance may hinder both firms effort to innovate and consumer pressure to do so. Consumers may internalize firms coordination problem. As a result, when improving environmental fitness of the good proves too difficult, firms may privilege other consumer preferences, such as price and use characteristics. We model and analyze these features in what follows.

3.8.1 Model Details

To model technological complexity related to the interaction among several product components, we introduce the interaction between the use characteristics of a good as a determinant of its environmental fitness. Borrowing from the literature on fitness landscapes (Kauffman and Levin, 1987), we use a continuous version of the NK model (Valente, 2014) and assume that firms need to improve the environmental fitness of a good on a complex technological landscape, which shape is given by the interaction between the characteristics.

As discussed in Section 3.3, the environmental fitness of a good ($s(\vec{x}_i)$) is given by the average fitness over its use characteristics (Eq. 15). For simplicity, we can assume that each characteristic is provided by one product component – for example the battery, RAM and CPU of a computer would determine, respectively, its off-the-grid autonomy, capability of dealing with several processes, and processing speed. However, unless the product is perfectly modular, the environmental fitness of each of these use characteristics *h* in a given paradigm *z* for a firm *i* ($\varphi_{z,i,h}$) depends on its own position on the technological landscape ($x_{i,h}$), as well as on the position of the other use characteristics that are part of the same technological paradigm *z*:

$$\varphi_{z,i,h} = \frac{\hat{s}(z)}{1+|x_{i,h}-v_{i,h}|} \tag{22}$$

where $\hat{s}(z)$ is the maximum environmental fitness achievable in paradigm *z*; and $v_{i,h}$ is a variable that measures the relation between different characteristics, which depend on the position of the other characteristics on the landscape ($x_{i,g\neq h}$) and on the strength of the relation between x_g and x_h ($a_{g,h}$):

$$v_{i,h} = \chi_{z,i} + \sum_{g \in z^i} a_{g,h} x_{i,g\neq h} \tag{23}$$

where $\chi_{z,i}$ is a variable that measures the fitness of each use characteristic with respect to it optimal position ($\hat{\chi}_z$), i.e. the value that firms attempt to attain to minimize the environmental impact of their good:

$$\chi_{z,i} = \hat{\chi}_z - \sum_{h \in z^i} \hat{\chi}_z a_{g,h} \qquad (24)$$

The crucial parameter here is $a_{g,h}$. When $a_{g,h} = 0$, $\chi_{z,i} = \hat{\chi}_z$ and $v_{i,h} = \chi_{z,i}$: it is sufficient to improve each characteristic towards their optimal position to reach the maximum environmental fitness. This is the model we discussed in the earlier Sections. When $a_{g,h} > 0$, each change in a characteristic determines the contribution of all others to the product environmental fitness. As $a_{g,h}$ approaches 1 the technological landscape is extremely complex because even small changes in the position of one characteristic has a strong impact on the relative fitness contribution of all other characteristics. Back to our computer example, improving the CPU will also reduce the off-the-grid autonomy for a given battery.

As explained in Section 3.4, each firm, in each time period, attempts to innovate. If successful, the firm has the opportunity to change the position of one characteristic *h* on the landscape (*mutation*). Differently from the above model, we now introduce the feature that a change in the position $x_{i,h}$ of characteristic *h* on the technological landscape changes the contribution to the environmental fitness also of characteristics $g \neq h$.

When moving to the second innovation step (*evaluation*) the firm considers if the utility of the target consumer group increases as a result of the innovation. The outcome now depends not only on the trade-off between quality of the good characteristics, price, and environmental fitness, but also on the technological complexity. That is, on whether the change of one characteristic towards the optimal position determines an overall environmental fitness improvement or not, even when its own fitness improves (the specific component that is changed is less polluting). As before, an increase in $x_{i,h}$ results in an increase in the direct utility, and a decrease in the indirect utility (Eq. 2) to an extent that depends on the consumer class preferences. But it is not foreseeable what happens to the environmental component of the utility, which may increase or not, depending on the complex interaction between the different product characteristics. Even more important, as shown in the literature on NK fitness landscapes, mutations that lead to increases in the environmental fitness may easily lead to a local optimum, that is a condition from which no other mutation would yield to an increase in the environmental fitness (unless the firm is able to make several steps at a time, along different product characteristics, and move to a different part of the technological landscape – something that we do not model here). As a result, firms that may not find ways to improve the environmental fitness, for example because they get stuck in a local optimum on the landscape, may decide to improve the quality of the product characteristics, which are appealing to the consumer preferences with respect to quality, and abandon efforts to improve the environmental fitness.

In what follows, we use this extended version of the model to study how the sheer complexity of the technology may influence the transition towards the production of greener goods.

3.8.2 Results

For the sake of comparability with the results discussed in Section 3.7, we use the same benchmark initialization of the model described in Appendix A, with equal consumer classes, all endowed with the same preferences, except for the product characteristics which are randomly distributed (with same support of the distribution). Differently from results discussed in Section 3.7, we now allow $a_{g,h}$ to change, to study the role of technological complexity and uncertainty on environmental outcomes.

Using the extended version of the model, the first property that we study is the relation between a measure of technological complexity $a_{g,h}$ and pollution stock. Figure 1 shows that as the technological complexity increases, as expected, the pollution stock in the final time step of the simulation is also higher. This is because for firms it is more difficult to improve the environmental fitness of their product, reach the peak of a technological paradigm, and therefore also move to new, less polluting, paradigms (See Figure 5 in Appendix B). Figure 1 also shows that the relation is not linear: small changes in complexity and uncertainty over innovation outcomes for low and high complex technologies have a small impact on the pollution stock. The largest impact on pollution occurs for intermediate values of technological complexity, suggesting that policy interventions are best targeted at supporting innovation or reducing uncertainty for products that are integrated, but not too much.

[Figure 1 about here]

To understand how this result emerges from firm behavior, and their interaction with the heterogenous consumer classes, Figure 2 plots the average distance across all product characteristics with respect to their optimal position in the technological landscape of a the technological paradigm on which they operate when we start the simulations. The optimal position is the one that would attain the maximum level of environmental fitness in a paradigm. In this configuration the optimal position in the initial paradigm is two, and the $x_{i,h}$ can range between 0 and 2.5 (they are initially assigned a position at random between the minimum and the maximum values on their space). So, a distance equal to 0.2 means that, the average product characteristic, across all firms, and over all simulations and simulated time steps stands at 1.8 or 2.2. This is high enough to find a balance between consumer preferences on the product characteristics and environmental fitness, without having to reach maximum environmental fitness. We report the average over all periods and across firms. The full fitted line is the simple average across firms, whereas the dashed fitted line is the average weighted by firm market shares.

Results in Figure 2 show that, in the absence of complexity ($a_{g,h} = 0$), there is a significant number of firms that compete without improving the environmental fitness of the good (the simple average distance of use characteristics with respect to optimal position is around 0.15). However, these firms tend to have a tiny market share: when weighted by market shares, the average distance with respect to the optimal position is nearly 0. This implies that the firms that succeed and grow must improve also the environmental fitness of the produced good.

Second, Figure 2 shows that as complexity increases (for larger values of $a_{g,h}$), on average firms perform worse with regards to environmental fitness. This is in line with the results discussed in Figure 1. As noted then, beyond middle levels of complexity $\left(a_{g,h} \succ 0.5\right)$ firm performance in terms of increasing environmental fitness within a paradigm does not worsen. This confirms that the impact of technological complexity on pollution stock (Figure 1) is due to firm inability to discover new technological paradigms.
Third, the results using the weighted average in Figure 2 suggest that the best performing firms, even in the presence of high complexity/uncertainty, manage to find close to optimal position for most of the product characteristics: the average distance is rather low, even for intermediate levels of complexity.
Fourth, we note that for the best performing firms (those with larger market shares) the distance even reduces as complexity moves beyond middle levels. This is because, as noted in Figure 5 in Appendix B, for very high levels of complexity, new paradigms are less likely to emerge. Therefore, new firms tend to enter in lower (on average more polluting) paradigms. The larger the number of firms that attempt to improve the environmental fitness in a technological paradigm, starting each time from a different random position, the larger the probability that at least one will find their way to the global maximum (those who start with product characteristics that are already very close to the maximum). In other words, in the presence of high technological complexity/uncertainty, firms are better off exploiting mature paradigms, than exploring new ones. This reduces the number of emerging, less polluting, technological paradigms, and contributes to increase the pollution stock, even if the most successful firms in mature paradigms, on average, produce more sustainable goods.

[Figure 2 about here]

The above results on firm environmental performance depend on consumer behavior and leverage. We show this in Figure 3, where we plot the minimum environmental fitness of goods tolerated by consumers, for varying level of the technological complexity. As discussed in Eq. 4, as pollution increases, the minimum level of environmental fitness tolerated increases logistically, until it reaches a given ceiling which represents the maximum fitness attainable within a given technological paradigm (which the average firm, as shown in Figure 2, cannot reach). The ceiling increases in the next technological paradigm, requiring firms to reach a higher fitness to compete in new paradigms.
Figure 3 shows the average minimum environmental fitness tolerated by consumers (for increasing levels of technological complexity) in three different time steps: after 250 periods (red crosses); after 380 time steps (green circles with crosses in the middle); and at the end of the simulation, after 3000 time steps (blue circles). Starting from the last period (blue hollowed circles), the results show that as complexity increase, the minimum level of environmental fitness that attain a positive environmental utility decrease with the level of technological complexity/uncertainty. This seems at odds with the result that, with higher complexity also total pollution increases (Fig. 1): with higher level of pollution, consumers

utility decreases drastically, which in turn decreases their tolerance with respect to pollution. This should also increase the minimum level of environmental fitness they demand. However, consumers form their expectation within a technological paradigm. Therefore, unless a new paradigm is explored, or comes to existence, consumers expectation are tailored on what can be reasonably expected from technological improvement within that existing mature paradigm. Because technological complexity reduces the pace at which new paradigms emerge, this also keeps consumer expectation low. This is also why firms are better off exploiting the current paradigm than exploring new ones (as discussed with reference to Figure 2).

We show this better as we look at results in earlier time steps. At the beginning of the simulation, around about 250 time steps (red crosses in Figure 2), when all firms are still exploring the existing paradigm (no new one has emerged), complexity increases the minimum level of environmental fitness tolerated in a good, because firms advance slowly on a complex landscape, causing a fast increase in pollution. After about 380 time steps complexity has no impact on the minimum level of environmental fitness tolerated in a good, because with low complexity new paradigms emerge, raising expectation (and therefore the minimum level), whereas with high levels of complexity new paradigms do not emerge but consumers react to the increasing pollution, increasing the minimum level. For low levels of complexity, when it is easy to improve the environmental fitness, the increase in the minimum level is driven by increased expectations about the technology, leading to a virtuous cycle and lower pollution; for high level of complexity the increase in the minimum level does not lead to more green innovation: consumer keep demanding for less polluting goods, but firm innovation alone cannot satisfy them.

[Figure 3 about here]

Because firms find it so difficult to improve the environmental fitness, in the presence of a complex technology, they have a stronger incentive to address different preferences to preserve higher market shares: use characteristics and price. We show this analyzing the relation between technological complexity and consumer utility. Figure 4 plots the relation between technological complexity and average consumer utility (panel a) and the weighted average of consumer preferences for product characteristics (using consumer classes population shares as weights). Because the environmental fitness falls with technological complexity, so does overall consumer utility, the more so, the less firms are able to address the negative externalities of increased pollution stock (Figure 4, panel a).

However, the decrease in overall utility is relatively low because firms can target other preferences, such as price and use characteristics, keeping consumers relatively happy and polluted. Because with high technological complexity steering the product characteristics towards a higher environmental fitness rarely works, firms tend to target classes with higher preferences for the product characteristics (recall that the quality of each product characteristic is independent from the quality of the other characteristics: in the model we assume that it is only the environmental fitness that depends on their interaction). This is what we show in Figure 4, panel b, where we plot the relation between technological complexity and the preferences of the surviving consumed classes for quality (rather than environmental fitness). In the presence of high technological complexity, the classes that thrive, and attract more consumers, are those with higher preferences for the quality of product characteristics. In other words, in the presence of high technological complexity a simple improvement of product characteristics within the same paradigm has higher returns, on average, than attempts to increase the environmental fitness. Also for this reason, firms that exploit the existing paradigm rather than exploring new more sustainable ones are able to thrive (as discussed in relation to Figure 2).

To sum up, the expectation of technological solutions to environmental problems may not be sufficient to reduce the environmental impact of consumption. This is especially the case in the presence of technological complexity, uncertainty with respect to the relation between innovation and environmental fitness, or lack of coordination between the producers of different component of a good along a value chain (see discussion in Section 4). In such situations, firms alone may easily find themselves locked in suboptimal areas of the technological landscapes, unable to improve the environmental fitness. Because of this, firms alone are also less likely to discover new paradigms when dealing with complex technologies. Slow paradigm shifts (exploitation of mature paradigms, rather than exploration of new ones) give firms an opportunity to reduce pollution only marginally, while improving product characteristics, and avoiding to move to a new, unknown, complex landscape. This is also because consumers pressure is less effective with high technological uncertainty (complexity/lack of coordination), as consumers tend to imitate peers that care more about the use characteristics than about the environmental fitness of the consumed good.

4. Discussion and Concluding Remarks

We model an economy in which improving the environmental fitness of goods produced by private firms is not straightforward. The interactions between different use characteristics of a product (e.g. the speed, security, autonomy of a car), as well as between the three populations of consumers, producers and products make any innovation process complex and uncertain. Here we are mainly interested in the environmental outcome of this process, i.e. whether the new good is more or less polluting (i.e. its environmental fitness is lower or higher).

[Figure 4 about here]

In the model, firms face trade-offs between three features of their products, which they can improve to satisfy consumers preferences, namely price, environmental fitness, and a vector of characteristics that define the good quality. For example, improving the environmental fitness may require to change the quality attached to the use characteristics of the product (e.g. the autonomy of an electric car). As a result of firms focusing on improving selected product characteristics that improve the utility of the consumers that they target, they may fail to see that along different innovation pathways substantially bigger improvements in the environmental fitness were possible. Such complexity may be difficult to govern by individual firms that control only parts of the production process, and therefore can innovate only on some of the product characteristics. To capture such complexity we model a scenario in which firms attempt to improve their product searching on a complex technological landscape. That is, a landscape that has multiple dimensions and which they discover as they explore it. All dimensions are interdependent, so changes in one dimension simultaneously affects all other dimensions, and the way in which they influence the environmental fitness of the product. As they explore the landscape, firms target evolving consumer preferences with respect to the quality of use characteristics, price, and environmental fitness (i.e. low pollution). Because of trade-offs between quality, price and environmental fitness and because of heterogeneous consumer preferences, depending on firm innovation some consumers benefit more than others and enjoy higher utility. In our model consumers imitate these better-off consumers, and slowly most consumers become mostly interested in the features that dominate the market (the quality of a given characteristic, price or environmental fitness). If most consumers benefit from the quality of a given product characteristic, because of the way in which the interactions between them the firms and the technology evolves, other consumers will also slowly become less interested in the environmental fitness of what they purchase. Because in the presence of complex technologies firms find it difficult to improve the environmental fitness, it is also less likely that new paradigms, with less polluting technologies, are discovered. As a result, because expectation with respect to environmental fitness vary across paradigms (as consumers can only compare technologies that exist), consumer expectations with respect to environmental fitness remain low, and firms remain stuck in mature, polluting, paradigms.

Technological complexity introduces a vicious cycle. Firms tend to remain stuck in local optima, and stop innovating towards more environmentally friendly goods. As a result, the likelihood of finding a new paradigm reduces, which also reduces the likelihood that consumers raise their expectations with respect to the environmental fitness. Because consumer expectations settle on low environmental fitness paradigms, firms are better off exploitation mature paradigms and improving the quality of product characteristics rather than exploring new paradigms, which would increase environmental fitness potential but which would also increase consumer expectations. As a result, the consumers that enjoy the highest utility are those that have lower preferences for environmental fitness. Firms then have an even stronger incentive to remain in the same technological paradigm and to increase the utility of consumers by improving the product characteristics that they prefer. Slow paradigm shift (exploitation) allows more firms to reduce pollution marginally, while at the same time improving product characteristics, rather than moving to a new, unknown, technological landscape.

Such vicious cycle requires a policy intervention, which is not modelled here, but would be an interesting extension to the toy model we proposed. How would public sector research influence the results, for example by exploring a wider space of the technological landscape, reducing the lock-in on optimal solutions? Regulations may also increase the coordination between producers of different component of final goods, for example by setting environmental standards.

The model proposed in this paper can be applied to study diffusion of green products in a number of industries. It is especially applicable to study diffusion of complex, non-modular technologies and products, the use of which is energy intensive, such as electric cars, refrigerators and air-conditioners. For instance, it could be used to study a radical transformation of the transport system, while promoting the diffusion of electric vehicles. The electrification of transport is expected to reduce CO2 emissions, ease gasoline-dependency, and improve cities' air quality. However, electric cars are in general perceived as less reliable than conventional vehicle because of their low battery lifespan. In addition, a high purchasing price of electric vehicles creates a barrier for many consumers. There are high hopes that environmentally-conscious consumers would be among early adopters, driving transitions to electric cars. In the example of our model this would require a large group of 'eco warriors' that are numerous and strong enough to attract more consumers. As we show in Windrum et al. (2009b) this also risks that a relatively small number of eco-warriors are well catered by green firms, and are enough for other consumers to free ride and consume goods with low price and environmental fitness.

Moreover, consumers are often unable to assess the overall environmental impacts of different car designs, as these depend on pollution created in the process of generating electricity, environmental impact of batteries and other technical features. As a result, consumers compare uncertain environmental performance against other car characteristics. This can make firms prefer to improve characteristics different from environmental performance, which in turn would make them fail to achieve maximum reductions in pollution. For instance, the emission reductions due to improvements in fuel-efficiency have been offset by increasing the average weight of new cars. Introducing environmental standards is one way to direct firms innovative effort to improve environmental performance of their car designs. In this context, the model proposed here can help to study optimal policies to guide innovations by firms towards improving environmental performance of their products and how this would affect a general technological progress in the automobile industry.

Another example of an application is e-waste in the computer industry. The fast advance in computational power (Moore's law) has been accompanied by a fast advance in the production of software. Hardware and software are strongly integrated, and their production is decentralized across different producers. Producers in both industries attempt to improve their profits, by improving the fitness of the use characteristics on which they specialize, for instance computer power (CPU) and software functionalities. As computer power increases, software producers have a lower incentive to increase the efficiency of the software, and a higher incentive to add features. Consumers also have different preferences and needs, and each benefit from the addition of functionalities in the software that may be used by a small portion of consumes. These consumers demand constantly more powerful computers, which causes existing computers to become quickly outdated, generating e-waste. Firms in the computer and software industry do not search to coordinate actions to reduce e-waste, and therefore the impact of computers on the environment. Instead, they focus on the innovation in product characteristics, to appeal to most consumers, who have little information on and rather uncertain exactions about the environmental impact of computers.

The model misses several relevant aspects that may allow to address the complexity and that suggest useful extensions for policy making. For instance, R&D has no cost in this version of the model, which may make firms incentives to move to a new paradigm even lower. Unless the demand or policy constraints are large enough. We encourage the use of the code in the modular LSD application[6] to extend the model in several useful directions and applications.

References

Alkemade, Floortje, Koen Frenken, Marko P. Hekkert, and Malte Schwoon, “A complex systems methodology to transition management,” Journal of Evolutionary Economics, aug 2009, 19 (4), 527–543.

Archibugi, Daniele. 2017. “Blade Runner Economics: Will Innovation Lead the Economic Recovery?” *Research Policy* 46 (3): 535–43.

Arthur, W Brian, “Competing technologies, increasing returns and lock-in by historical events,” Economic Journal, 1989, 99, 116–131.

Auerswald, Philip, Stuart Kauffman, Jos´e Lobo, and Karl Shell, “The production recipes approach to modeling technological innovation: An application to learning by doing,” Journal of Economic Dynamics and Control, mar 2000, 24 (3), 389–450.

Balint, T., F. Lamperti, A. Mandel, M. Napoletano, A. Roventini, and A. Sapio, “Complexity and the Economics of Climate Change: A Survey and a Look Forward,” Ecological Economics, aug 2017, 138, 252–265.

Barker, Terry, Athanasios Dagoumas, and Jonathan Rubin, “The macroeconomic rebound effect and the world economy,” Energy Efficiency, may 2009, 2 (4), 411–427.

Bleda, Mercedes and Marco Valente, “Graded eco-labels: A demand-oriented approach to reduce pollution,” Technological Forecasting and Social Change, 2009, 76 (4), 512–524.

Buenstorf, Guido and Christian Cordes, “Can sustainable consumption be learned? A model of cultural evolution,” Ecological Economics, nov 2008, 67 (4), 646–657.

Ciarli, Tommaso, and Maria Savona. 2019. “Modelling the Evolution of Economic Structure and Climate Change: A Review.” *Ecological Economics* 158 (April): 51–64.

_ , Andr´e Lorentz, Marco Valente, and Maria Savona, “Structural Changes and Growth Regimes,” Journal of Evolutionary Economics, 2018, Online.

_ , Riccardo Leoncini, Sandro Montresor, and Marco Valente, “Technological change and the vertical organisation of industries,” Journal of Evolutionary Economics, 2008, forthcomin.

Cincotti, Silvano, Marco Raberto, and Andrea Teglio, “Credit Money and Macroeconomic Instability in the Agent-based Model and Simulator Eurace,” Economics: The OpenAccess, Open-Assessment E-Journal, 2010, 4 (2010-26), 1.

Cowan, Robin and Philip Gunby, “Sprayed to Death: Path Dependence, Lock-in and Pest Control Strategies,” The Economic Journal, may 1996, 106 (436), 521–542.

Dosi, Giovanni, “Technological paradigms and technological trajectories : A suggested interpretation of the determinants and directions of technical change,” Research Policy, 1982, 11(3), 147–162.

_ , Giorgio Fagiolo, and Andrea Roventini, “Schumpeter Meeting Keynes: A Policy– Friendly Model of Endogenous Growth and Business Cycles,” Journal of Economic Dynamics and Control, 2010, 34 (9), 1748–1767.

Farmer, J. Doyne, Cameron Hepburn, Penny Mealy, and Alexander Teytelboym, “A Third Wave in the Economics of Climate Change,” Environmental and Resource Economics, oct 2015, 62 (2), 329–357.

Frenken, Koen, Luigi Marengo, and Marco Valente, “Interdependencies, {N}early{D}ecomposability and {A}daptation,” in Thomas Brenner, ed., Computational Techniques for Modelling Learning in Economics, Boston Dordrecht and London: Kluwer, 1999.

Gallouj, Faiz and Olivier Weinstein, “Innovation in Services,” Research Policy, 1997, 26, 537–556.

Gerst, M.D., P. Wang, A. Roventini, G. Fagiolo, G. Dosi, R.B. Howarth, and M.E. Borsuk, “Agent-based modeling of climate policy: An introduction to the ENGAGE multilevel model framework,” Environmental Modelling & Software, jun 2013, 44, 62–75.

Grin, John, Jan Rotmans, and Johan Schot, “Transitions to Sustainable Development. New Directions in the Study of Long Term Transformative Change,” 2010.

Hafner, S., Anger-Kraavi, A., Monasterolo, I., Jones, A., (2020) Emergence of new economics energy transition models: a review. Ecological Economics 177, forthcoming.

Janssen, Marco A. and Wander Jager, “Stimulating diffusion of green products,” Journal of Evolutionary Economics, jul 2002, 12 (3), 283–306.

Johnstone, Phil and Andy Stirling, “Comparing Nuclear Power Trajectories in Germany And The UK: From Regimes’ to Democracies’ in Sociotechnical Transitions and Discontinuities,” 2015.

Kauffman, Stuart A., The origins of order : self-organization and selection in evolution, Oxford University Press, 1993.

Kauffman, Stuart A and Simon Levin, “Towards a general theory of adaptive walks on rugged landscapes,” Journal of Theoretical Biology, 1987, 128 (1), 11–45.

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

Klepper, Steven. 1996. “Entry, Exit, Growth, and Innovation over the Product Life Cycle.” *American Economic Review* 86 (3): 562–83.

Lamperti, F., G. Dosi, M. Napoletano, A. Roventini, and A. Sapio, “Faraway, So Close: Coupled Climate and Economic Dynamics in an Agent-based Integrated Assessment Model,” Ecological Economics, aug 2018, 150, 315–339.

Lamperti, Francesco, Antoine Mandel, Mauro Napoletano, Alessandro Sapio, Andrea Roventini, Tomas Balint, and Igor Khorenzhenko. 2019. “Towards Agent-Based Integrated Assessment Models: Examples, Challenges, and Future Developments.” *Regional Environmental Change* 19 (3): 747–62.

Lancaster, K J, “A New Approach to Consumer Theory,” Journal of Political Economy, 1966, 14, 133–156.

_ , “Change and Innovation in the Technology of Consumption,” The American Economic Review, 1966, 56(1/2), 14–23.

Langlois, Richard N., “Modularity in technology and organization,” Journal of Economic Behavior & Organization, sep 2002, 49 (1), 19–37.

Leach, Melissa, Ian Scoones, and Andrew Stirling, “Pathways to Sustainability: an overview of the STEPS Centre approach,” 2007.

Malerba, Franco and Luigi Orsenigo, “Technological Regimes and Sectoral Patterns of Innovative Activities,” Industrial and Corporate Change, 1997, 6 (1), 83–118.

Merton, Robert C, “Optimum consumption and portfolio rules in a continuous-time model,” Journal of Economic Theory, dec 1971, 3 (4), 373–413.

Mowery, David and Nathan Rosenberg, “The influence of market demand upon innovation: a critical review of some recent empirical studies,” Research Policy, 1979, 8 (2), 102–153.

Nelson, Richard Robert and Sidney G Winter, An Evolutionary Theory of Economic Change, Cambridge, MA: Harvard University Press, 1982.

Oltra, Vanessa and Ma¨ıder Saint Jean, “Sectoral systems of environmental innovation: An application to the French automotive industry,” Technological Forecasting and Social Change, may 2009, 76 (4), 567–583.

Pasimeni, Francesco and Tommaso Ciarli, “Diffusion of Shared Goods in Consumer Coalitions. An Agent-Based Model,” 2018.

Pindyck, Robert S, “Climate Change Policy: What Do the Models Tell Us?,” Journal of Economic Literature, sep 2013, 51 (3), 860–872.

Ponta, Linda, Marco Raberto, Andrea Teglio, and Silvano Cincotti, “An Agent-based Stock-flow Consistent Model of the Sustainable Transition in the Energy Sector,” Ecological Economics, mar 2018, 145, 274–300.

Safarzynska, Karolina and Jeroen C. J. M. van den Bergh, “An evolutionary model of energy transitions with interactive innovation-selection dynamics,” Journal of Evolutionary Economics, apr 2013, 23 (2), 271–293.

_ and Jeroen C.J.M. van den Bergh, "Evolving power and environmental policy: Explaining institutional change with group selection," Ecological Economics, feb 2010, 69 (4), 743–752.

_ , Koen Frenken, and Jeroen C.J.M. van den Bergh, "Evolutionary theorizing and modeling of sustainability transitions," Research Policy, jul 2012, 41 (6), 1011–1024.

Saviotti, Pier Paolo and J Stan Metcalfe, "A Theoretical Approach to the Construction of Technological Output Indicators," Research Policy, 1984, 13, 115–141.

Scoones, Ian, Melissa Leach, Adrian Smith, Sigrid Stagl, Andy Stirling, and John Thompson, "Dynamic Systems and the challenge of Sustainability," 2007.

Simon, Herbert A, "Near decomposability and the speed of evolution," Industrial and Corporate Change, 2002, 11 (3), 587–599.

Stapleton, Lee, Steve Sorrell, and Tim Schwanen, "Estimating Direct Rebound Effects for Personal Automotive Travel in Great Britain," SPRU Working Paper Series, 2015, (08).

Stern, Nicholas, "The Structure of Economic Modeling of the Potential Impacts of Climate Change: Grafting Gross Underestimation of Risk onto Already Narrow Science Models," Journal of Economic Literature, sep 2013, 51 (3), 838–859.

_ , "Economics: Current climate models are grossly misleading," Nature, feb 2016, 530 (7591), 407–409.

Steve Sorrell, "The Rebound Effect: an assessment of the evidence for economy-wide energy savings from improved energy efficiency," UK Energy Research Centre, UERC 2007.

Tedeschi, Gabriele, Amin Mazloumian, Mauro Gallegati, and Dirk Helbing, "Bankruptcy Cascades in Interbank Markets," PLoS ONE, dec 2012, 7 (12), e52749.

Tesfatsion, Leigh, "Agent-Based Computational Economics: A Constructive Approach to Economic Theory," in L Tesfatsion and K L Judd, eds., Agent-Based Computational Economics, Vol. 2 of Handbook of Computational Economics, Elsevier, 2006, chapter 16, pp. 831–880.

Thurner, Stefan and Sebastian Poledna, "DebtRank-transparency: Controlling systemic risk in financial networks," Scientific Reports, dec 2013, 3 (1), 1888.

Tsur, Yacov and Amos Zemel, "Towards endogenous recombinant growth," Journal of Economic Dynamics and Control, nov 2007, 31 (11), 3459–3477.

Valente, Marco, "An NK-like model for complexity," Journal of Evolutionary Economics, oct 2014, 24 (1), 107–134.

van den Bergh, Jeroen C.J.M., Bernhard Truffer, and Giorgos Kallis, "Environmental innovation and societal transitions: Introduction and overview," Environmental Innovation and Societal Transitions, jun 2011, 1 (1), 1–23.

Weitzman, M. L., "Recombinant Growth," The Quarterly Journal of Economics, may 1998, 113 (2), 331–360.

Windrum, Paul and Chris Birchenhall, "Is product life cycle theory a special case? Dominant designs and the emergence of market niches through coevolutionary-learning," Structural Change and Economic Dynamics, 1998, 9 (97), 109–134.

_ , Tommaso Ciarli, and Chris Birchenhall, "Consumer heterogeneity and the development of environmentally friendly technologies," Technological Forecasting and Social Change, 2009, 76 (4), 533–551.

_ , _ , and _ , "Environmental impact, quality, and price: Consumer trade-offs and the development of environmentally friendly technologies," Technological Forecasting and Social Change, 2009, 76 (4), 552–566.

Wirkierman, Ariel L., Tommaso Ciarli, and Mariana Mazzucato, "An evolutionary agent-based model of innovation and the risk-reward nexus," Technical Report, SPRU, University of Sussex, Brighton, UK 2018.

Wolf, Sarah, Steffen Fürst, Antoine Mandel, Wiebke Lass, Daniel Lincke, Federico Pablo-Martí, and Carlo Jaeger, "A multi-agent model of several economic regions," Environmental Modelling & Software, jun 2013, 44, 25–43.

Zeppini, Paolo and Jeroen C. J. M. van den Bergh, "Competing Recombinant Technologies for Environmental Innovation: Extending Arthur's Model of Lock-In," Industry & Innovation, apr 2011, 18 (3), 317–334.

A Initialization

We set up a benchmark configuration with average values of the critical parameters (Table 1). Consumers preference toward the environmental sustainability of goods (η) is fixed and equal across consumer classes; similarly for indirect preference (α); direct preferences toward each product characteristic (β_h) are randomly drawn from a uniform distribution that is also equal across classes. In sum, benchmark results are an outcome of a random selection between consumer classes, which occurs as their preferences randomly change through time – as classes that enter the market bring novelty in consumption tastes; rather than an outcome of a selection on environmental preferences.

We run simulations with a population of 25 firms and 500 consumers divided into 100 consumer classes, all fixed through time. Both firms and classes start with equal endowments and equal share of sales and consumers respectively. We run each setting for 3000 time periods. Unless differently stated in the text, all result present average simulation outcomes over 10 different runs: after a preliminary analysis of the model we have considered this a good trade-off between results verification and computational effort. The interested reader may refer to Windrum et al. (2009a) for a sensitivity analysis.

Table 1: Parameters setting

Par / Var_{t-1}	**Description**	**Value**
$\hat{C}$	total number of consumers in the economy	500
f	replicator tamed parameter	5
ε^u	minimum survival term	0.02
m_j	Endowment	10
α_j	Indirect utility preference	0.5
$\beta_{j,h}$	Preferences for product characteristics	U [0.1,0.3]
η_j	Preference for environmental sustainability (discount rate)	0.6
ρ	Relative risk aversion toward pollution	0.5
η_j^p	Preference toward the actual environmental sustainability of the good rather than the potential sustainability of the techno–environmental paradigm	0.5
s_0	Lower asymptote of the minimum environmental fitness logistic function	0.01
r	Rate of growth of the minimum environmental fitness logistic function	5e-005
λ^y	Rate of adjustment of production decisions	0.2
λ^c	Rate of adjustment of capital stock to production needs	0.2

ν	Mark–up	0.1
$\hat{\zeta}$	Maximum environmental impact of goods	1
φ	Speed of impact reduction of a fitness increase	2
ι	Probability of success of innovation on one characteristic	0.2
ξ	Mutation weight	0.2
σ^Z	Peak variance that allows to open a new window of opportunity	0.005
τ^Z_{Min}	Minimum number of periods needed to discover a new techno-environmental paradigm	100
τ^Z_{Max}	Maximum number of periods needed to discover a new techno-environmental paradigm	50
Δ^{φ}	Change in the maximum level of environmental fitness across paradigms	0.5
$\boldsymbol{\sigma^{\chi}}$	Variance of the technological change of the environmental landscape	0.3
$\boldsymbol{\delta^c}$	Minimum number of consumers below which a class is replaced	2
$\boldsymbol{\delta^k}$	Minimum amount of capital below which a firm is replaced	0.2
$\boldsymbol{\tau^r_{Min}}$	Minimum number of periods between two firms and consumers turnovers	10
$\boldsymbol{\tau^r_{Max}}$	Maximum number of periods between two firms and consumers turnovers	20
$\boldsymbol{\beta_{Min}}$	Minimum value consumer preferences toward product characteristics	0.1
$\boldsymbol{\beta_{Max}}$	Maximum value consumer preferences toward product characteristics	0.3
$\boldsymbol{H_z}$	Number of user characteristics in any design in any paradigm	3
$\boldsymbol{x_{Min}}$	Minimum value of a product characteristic in the first period	0.1
$\boldsymbol{x_{Max}}$	Maximum value of a product characteristic in the first period	2.5
	Technological complexity	
$\boldsymbol{a_{g,h}}$	Environmental fitness interaction term: the effect of a change in x_g on the fitness of x_h	$U\ [a_{Min}, aa_{Max}]$
$\boldsymbol{a_{Min}}$	Minimum value of the product characteristics environmental fitness interaction	tested
$\boldsymbol{a_{Max}}$	Maximum value of the product characteristics environmental fitness interaction	tested

	Initial values	
τ_0^z	Initial number of periods to discover a new paradigm	50
$x_{0,i,h}$	Initial value of each product characteristic	$U[x_{Min}, x_{Max}]$
$\beta_{0,j,h}$	Initial value of the preference for each product characteristic	$U[\beta_{Min}, \beta_{Max}]$
z_0	Initial paradigm	1
τ_0^r	Initial number of periods before a firms and consumers replacement occurs	20
G_0	Initial level of environmental pollution	0
$\overline{u}_{0,j}^f$	Initial average (fuzzy) utility	1
$\hat{\chi}_{z=1}$	Optimal position of the x_h on the environmental landscape in the initial paradigm	2
ψ_0	Initial ratio of consumers per class	1/100

Figures

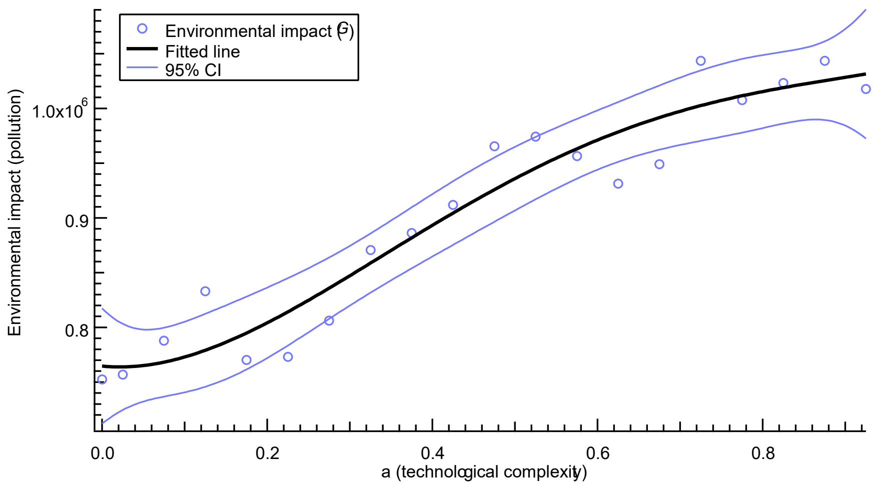

Notes. Fitted polynomial regression between different initial values of $a_{g,h}$ and average pollution stock in the last time period of the simulation.

Figure 1: *Relation between complexity ($a_{g,h}$) and environmental impact (G)*

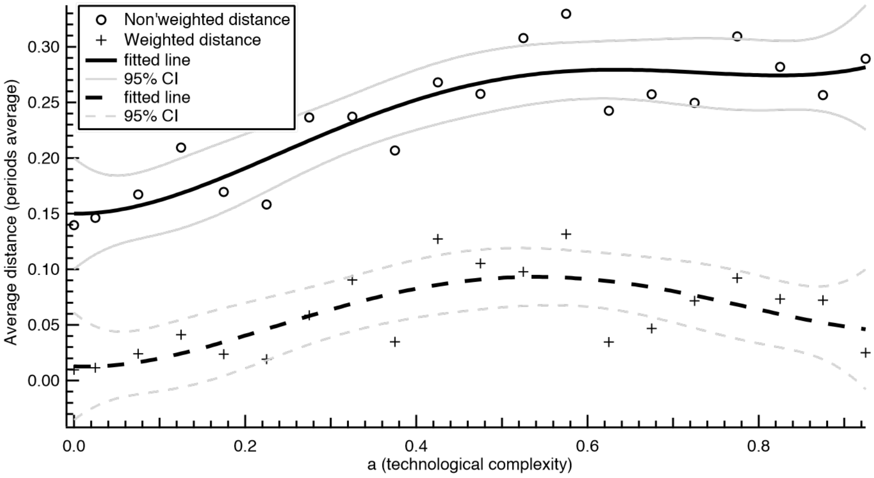

Notes. Fitted polynomial regression between different initial values of $a_{g,h}$ and the average distance between each characteristic and their optimal position (the one that attains maximum environmental fitness) across characteristics and firms. The average is further averaged across the 3000 time periods. The full fitted line (and hollowed circles) represents the simple average; the dash fitted line (and crosses) is the weighted average, using firm market share as weights

Figure 2: *Relation between complexity ($a_{g,h}$) and the average distance of product characteristics with respect to their optimal position*

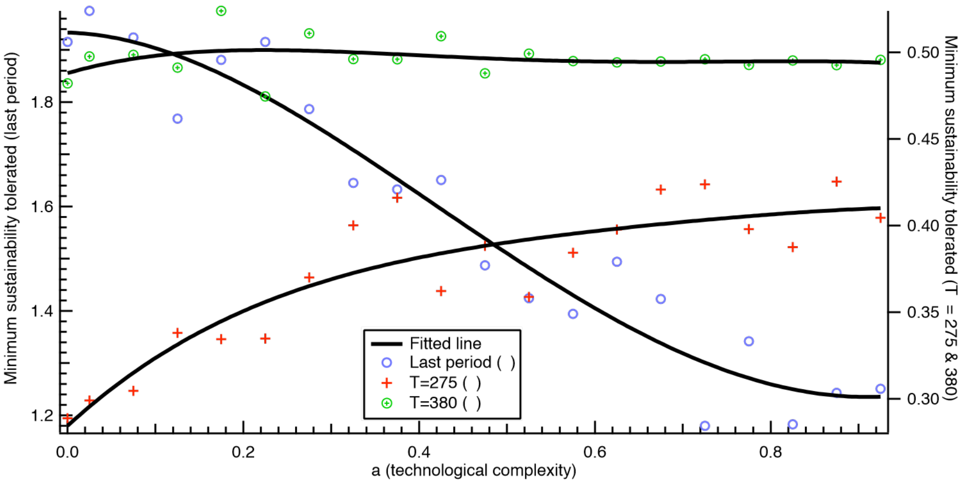

Notes. Fitted polynomial regression between different initial values of $a_{g,h}$ and the minimum level of good's environmental fitness accepted by consumers ($\underline{s}$) at the end of given periods. The red crosses are used to plot the relation after 250 periods; the green circles with crosses are used to plot the relation after 380 periods; the blue hollowed circles are used to plot the relation at the end of the simulation (3000).

Figure 3: *Relation between complexity ($a_{g,h}$) and the minimum level of good sustainability accepted by consumers*

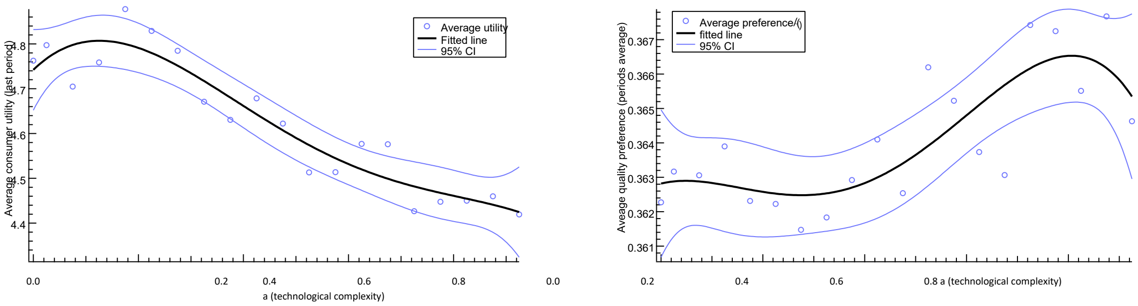

(**a**) Average Utility (**b**) Surviving classes' quality preferences

Notes. Fitted polynomial regression between different initial values of $a_{g,h}$ and average consumer utility in the last period (panel a) and the weighted average of the preferences with respect to product characteristics (β_h) across classes and periods.

Figure 4: *Relation between complexity ($a_{g,h}$) and consumer utility and preferences*

B Extra Figures

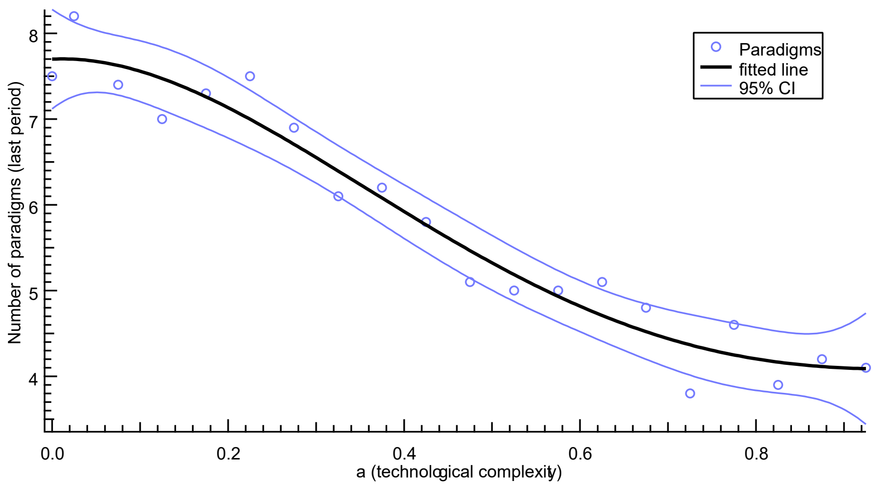

Notes. Fitted polynomial regression between different initial values of $a_{g,h}$ and average number of paradigm discovered in the last time period of the simulation.

Figure 5: *Relation between complexity ($a_{g,h}$) and the number of paradigms discovered*

[1] Paradigms *z* evolve endogenously in the model and are assumed to improve exogenously as an outcome of basic research – see section 3.5.

[2] In the unlikely event that two firms attain the same utility for a specific class, one of the two is randomly picked with equal probability.

[3] A modification of a truncated Cauchy function.

[4] Note that while the first term is strictly positive, the second can be negative: when a firm reaches the proximity of the maximum environmental fitness, it might not see the opportunity of a new paradigmatic search.

[5] As mentioned in section 3.5 at least one characteristic is in common to two consecutive paradigms.

[6] Laboratory for Simulation Development, available for free at the following website https://github.com/ marcov64/Lsd