Citation for published version:

Marengo, L & Zeppini, P 2016, 'The arrival of the new', *Journal of Evolutionary Economics*, vol. 26, no. 1, pp. 171-194. <https://doi.org/10.1007/s00191-015-0438-0>

DOI:

[10.1007/s00191-015-0438-0](https://doi.org/10.1007/s00191-015-0438-0)

Publication date:

2016

Document Version

Peer reviewed version

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

The arrival of the new

Luigi Marengo^a and Paolo Zeppini^b

^a *Department of Management, LUISS University, Roma, lmarengo@luiss.it*

^b *Department of Economics, University of Bath, p.zeppini@bath.ac.uk*

Abstract

In this work we present a number of urn models in which, contrary to standard Pólya urns, the number of competing alternatives is not given from the outset but may increase with the arrival of innovations. We begin by describing a variant of Pólya urns, first introduced by Fred Hoppe, in which balls of previously non-existing colors are added with some (declining) probability. We then propose new variants in which the probability of the arrival of new colors is itself subject to adaptive change depending on the success of past innovations and discuss applications to evolutionary models of technologies and industries. We numerically simulate different specifications of these urns with adaptively changing mutation rate and show that they can account for complex patterns of evolution in which periods of exploration and innovation are followed by periods in which the dynamics of the system is driven by selection among a stable set of alternatives.

1 Introduction

One of the most challenging questions in the social sciences is how an innovation (an innovative technology, a new business model, a new behavior or norm, etc.) can lead to a major transition in society. When innovations arrive, they compete with incumbents and can either gain a dominant position or coexist with them on a more or less equal basis or, on the contrary, fail and disappear. In turn, successful innovations will be challenged by the arrival of other novelties. The evolution of industries, markets and societies is in some sense the outcome of this complex interaction dynamics among a process of arrival of novelties and a process of competition among incumbent alternatives.

Models of path-dependence and lock-in have focused only on the latter and have been successful in explaining the probabilistic laws governing the competing diffusion process when a given set of options is available. Such models are often based on the so-called Pólya-urn schemes, special types of Markov chains where the probability distribution of outcomes depends on past draws. They have provided valuable insight in the study of path dependency in technological diffusion (Arthur, 1989; Arthur et al., 1987; Dosi et al., 1994), of reinforcement learning in signalling games (Erev and Roth, 1998), and in other economic models. However, these are “closed world” models that only analyze the stochastic processes of selection among a given set of options, for instance two or more competing technologies, under the implicit assumption that they come to existence at the same time and are not going to be challenged by further

innovations. Pólya-urn models are instead agnostic regarding the origin of such options and fail to address the interplay between the complex interaction between the processes of selection (adoption) and innovation. The result is that reinforcement mechanisms generate early lock-in and bring evolution to a stalemate.

An extension of Pólya-urn models has been proposed in theoretical biology by Hoppe (1984) as a benchmark model for neutral drift in evolution, where mutations do not carry any selective advantage. Hoppe-Pólya urns extend the standard Pólya-urns model by adding a stochastic process of addition of brand new options that probabilistically enter the competition for selection. Thus, potentially, Hoppe-Pólya urns can account for genuine innovation and persistence of variety. However, we will show, also in this model, the rate of innovation approaches zero rather quickly, and an initial phase of exploration is followed by endless exploitation. In this article, we propose a generalization of Hoppe-Pólya urn models where the rate of innovation is itself the outcome of some endogenous stochastic process.

In standard Pólya-urns, alternative options are represented as colored balls in an urn. At each time, one ball is drawn at random from the urn and reintroduced with an additional ball of the same color. Mutation is introduced in Hoppe-Pólya urns by adding a black ball that behaves as the other colored balls but acts as a *mutator*. If the black ball is drawn, a new color, which did not exist before, is added to the urn.

In this paper, we propose extensions of the Hoppe-Pólya model that, in our view, could cast new light on the complex interaction between the processes of selection among existing alternatives and of the generation of new alternatives. We will proceed by a series of small steps towards the development of the more interesting and comprehensive model in section 6. In a first step, we assume that the mutator itself is subject to self-reinforcement and we preserve neutrality. Each time the black ball is drawn in addition to adding a ball of a brand new color, another black ball is added, thus increasing the probability of other mutations happening.

In the second extension, we make the addition of the black ball conditional upon the success of the last innovation. Each color is attributed a random fitness value drawn from a given distribution. When the black ball is drawn, a new color is added: if its fitness value is higher than the average fitness of the colors in the urn the last innovation is considered successful and a new black ball is added. However, if the new fitness of the new color is below the average the number of black balls remains constant. The ball extraction process instead does not depend upon the fitness of the colors: as in standard Pólya urns, every ball in the urn has the same probability of being drawn.

Finally, in the third extension, we present what we believe is the key contribution of the paper, i.e. a fully-fledged model of interaction between selection and novelty generation in which we make the color extraction probability dependent upon the relative fitness of colors: fitter colors have a higher probability of being drawn and therefore being reinforced with the addition of a ball of the same color.

We provide a detailed study of the numerical properties of such variants of Hoppe-Pólya urn model with adaptive mutation rates. The distribution of simulation outcomes shows more variety than a simple Hoppe-Pólya model, with sustained novelty and different phases in the interplay between innovation and selection. The time evolution of the process shows that variety is not constant, but displays a pattern of punctuated equilibria. Thanks to the endogenous rate of mutation, successful innovations establish themselves, but after some time, newer successful innovations come about and take over the incumbents. The intuition is that while self-reinforcement of each option is marginally decreasing, the reinforcement mechanisms of the mutator is not, so, in the long run, the process of novelty is sustained.

The time required for a transition to occur is a function of the competition process between alternative innovations. Once a “winning” innovation has emerged from this competition, the incumbent’s decline begins (*“sic transit gloria mundi”*). The pattern of variety produced by our model may describe conceptually technological transitions. One example is the energy

sector, where a number of alternative renewable energy solutions (solar, wind, biomass, fuel-cells) compete to become the dominant alternative to fossil fuels. In a broader sense, our model can be relevant to more general societal transitions, proposing a probabilistic explanation of the rise and fall of organizations and societies.

To our knowledge, the only existing application of Hoppe-Pólya to a socio-economic problem is contained in Alexander et al. (2012). They use Hoppe-Pólya urns to analyze a signalling game (Crawford and Sobel, 1982) with reinforcement learning *à la* Erev and Roth (1998). They introduce the possibility that new signals can be stochastically introduced, which, in this context, becomes a way to escape pooling equilibria and increase the probability of convergence to separating equilibria. Once equipped with a reinforcement learning mechanism, Hoppe-Pólya urns are no longer neutral, but display selective advantage. In our variant, selective advantage operates not on the colored balls, but directly on the black one, therefore modifying the rate of mutation. This prevents early lock-in of the system in a limited set of options¹ and can produce sustained innovation.

The model proposed in this article may appear at a first sight as a rather abstract mathematical exercise and the purpose of this paper is mainly to propose what we think is a very interesting and promising tool for future analysis, rather than discussing any direct empirical application. Nevertheless, we believe the methodology we propose has some important relation to an emerging body of empirical studies. In a recent paper, Youn et al. (2014) study the secular record of patented innovations in the US since 1790 and the emergence of new patent classes, or new “technology codes”. Patents can either be attributed to existing codes, and therefore signal the reinforcement of an existing technological trajectory, or give rise to a novel code, i.e. a new technological class that signals a novel body of technological knowledge. Indeed, our model could be used precisely to give an analytic account of the emergence and decline of technological trajectories using this kind of patent data.

Similarly, Tria et al. (2014) present a framework somehow similar to ours in order to model the emergence and diffusion of novelties and apply it to four empirical cases: the sequence of words in a text (where novelty is the appearance of a new word), the edit events in Wikipedia (where novelty is a new edit event by a contributor who has never acted before), tagging events in a social annotation system and the access to different songs in an online music catalogue. The central claim of the paper is that the dynamics of novelties in these different domains seems to confirm the principle of the “adjacent possible” (Kauffman, 1993), which states that every novelty opens a new space of possibilities for additional novelties by making accessible a set of further novelties that lie “one step” away from it. The resulting process is one of ever increasing opportunities in an open world dynamics, and is in sharp contrast with the close world of standard Pólya urn models.

The principle of adjacent possible assumes a network (most typically a tree) structure of the underlying search space, according to which every novelty that is introduced makes it possible to access also all novelties that have a direct connection to it. A similar point is made by branching models of firm diversification (cf., for instance, Bottazzi and Secchi (2006)), which analyze the allocation process of a firm’s resources either to the growth of an existing product line or to the opening of a new one.

Here, we make a complementary point: we abstract from any spatial structure, but nevertheless show that, under some model specification, increasing returns do not necessarily lead to stable lock-in, but the arrival of novelties can always overturn the established incumbents.

The paper is organized as follows. In section 2, we briefly describe Hoppe-Pólya urn models and our variant with adaptive mutation rates. In section 3, we present, as benchmark cases, numerical results from the simulation of both standard Hoppe-Pólya urns and urns with non-adaptively increasing mutation rates, while, in sections 4 and 5, we describe the behavior of our

¹Alexander et al. (2012) model a very simple signalling game with a very small set of strategies and signals. In a larger game with more strategies and more signals, their results would be more problematic.

adaptive version when fitness values are drawn, respectively, from a uniform distribution and from a beta distribution. In section 6, we use fitness values not only to modify the probability of the introduction of new colors but also the probability of addition of balls of the existing colors. Finally, in section 7, we summarize and conclude.

2 The model

2.1 Hoppe-Pólya urns

Proposed by Fred Hoppe in *Journal of Mathematical Biology* in 1984 as a benchmark model for neutrality in evolution, Hoppe-Pólya urns model neutral selection, especially in those case in which there exist a vast number of potential mutations that convey no selective advantage, and extend Bayesian theory of induction, generalizing the Bayes-Laplace rules of succession allowing for the formation of new categories (Zabell, 1992).

A Hoppe-Pólya urn at time 0 contains only one black (“uncolored”) ball and at time n it contains 1 black ball and n colored balls of k different colors. At each time $n + 1$ one ball is randomly drawn from the urn. If a colored ball is drawn, it is reintroduced in the urn together with a new ball of the same color (as in a standard Pólya urn), thus at time $n + 1$ we have $n + 1$ colored balls of k different colors. (We remind once more that the black ball is not counted among the colored ones.)

If instead the black ball is drawn, it is reintroduced in the urn together with a new ball of a **new color**, and so at time $n + 1$ we have $n + 1$ colored balls of $k + 1$ different colors. More generally, we can suppose that the black ball has mass $\theta > 0$, whereas colored balls have all mass 1.

Obviously, in the urn process we just described, the number of colors grows logarithmically to infinity, as the expected number of colors at time n is given by the harmonic series $\sum_{i=1}^n 1/i$. Figure 1 plots the expected number of colors in the logarithm of time.

Figure 1: Expected number of distinct colors

A Hoppe-Pólya stochastic process can be formalized as follows. Let us call $\{1, 2, \dots\}$ the set of colors (once more, black is not considered a color, it is a mere “mutator”) and X_n the random variable that indicates the color of the new ball added to the urn after the n -th drawing. If we start with only one black ball in the urn, we have $X_1 = 1$, $X_2 = 1$ or 2 , $X_3 = 1, 2$ or 3 , etc..

Let k be the random number of distinct colors that are in the urn at time n , then the urn can be described as a set of k integers $\{n_1, n_2, \dots, n_k\}$ that form a partition² of the integer

²A partition of a positive integer n is a way of writing n as a sum of positive integers. For instance, the integer 4 can be partitioned in five ways: 4, 3+1, 2+2, 2+1+1, and 1+1+1+1.

n . Which one among the possible partitions of n will be realized depends upon the sequence $\{X_j\}_{j=1}^n$ of draws.

For each $1 \leq i \leq n$, call a_i the number of times the integer i appears in $\{n_1, n_2, \dots, n_k\}$, the vector $a = (a_1, a_2, \dots, a_n)$ is called the (allelic) partition vector. For instance, $a = (1, 0, 1, 0)$ indicates that there is one color with three balls and one color with one ball (thus $n = 4$).

For a given number of colors, the distribution of trials among the possible partitions is not uniform. Suppose, for instance, we have two colors after 10 trials (this happens with probability $\simeq 0.28$). This can be realized in 5 different ways as a partition of 10: $5+5$, $4+6$, $3+7$, $2+8$, $1+9$. As we will show below, the probability of each partition is given by Ewens' sampling formula (Ewens, 1972):

$$P[a_1, \dots, a_n] = \prod_{j=1}^n \frac{1}{j^{a_j} a_j!} \quad (1)$$

As an example, Figure 2 plots the probabilities of each such partitions for the case of 10 trials and 2 colors. It can be noticed that more “unequal” partitions are more likely. This property is called *preferential attachment*, meaning that more crowded cells are more likely to attract new tokens, and obviously reflects the property of increasing returns of the urn process.

Figure 2: Probabilities of partitions of 10 trials into 2 colors

Assuming that the initial mass of the black ball is θ (normally θ is set equal to 1), the fundamental result in Hoppe-Pólya urns is the following:

Hoppe's Theorem: the random partition Π_n determined by the sequence $\{X_j\}_{j=1}^n$ is a Markov process with marginal distribution:

$$P[\Pi_n = a] = \frac{n!}{[\theta]^n} \prod_{i=1}^n \frac{\theta^{a_i}}{i^{a_i} a_i!} \quad (2)$$

where $[\theta]^n = \theta(\theta + 1) \dots (\theta + n - 1)$.

The proof, based upon combinatorial arguments, can be found in Hoppe (1984).

2.2 Urns with adaptive mutation rates

In Hoppe-Pólya urns, the rate of mutation quickly vanishes as the probability of drawing the black ball at time n is $\theta/(\theta + n - 1)$. Thus we tend to observe early lock-in and the disappearance of any further exploration. This is especially unfortunate if the state space is large and/or non-stationary.

We propose, therefore, a variation of Hoppe-Pólya urns in which the mutation rate is itself subject to variability and adjusts in an adaptive fashion. First of all, we have to depart from the assumption of neutrality and suppose that colors have differential adaptation to the environment expressed by a fitness value. When a new color is created, its fitness is randomly drawn from some given distribution (we will experiment both with a uniform and a beta distribution). If the new color has a fitness value higher than the average fitness in the population of existing colored balls, then the mutation is considered successful and **a new black ball** is also added. If, instead, the new color has lower than average fitness, the mutation is considered unsuccessful and the number of black balls is kept unchanged.

To summarize, suppose that after $n - 1$ draws we have $n - 1$ colored balls of k different colors and b black balls, then at time n :

- if a colored ball is drawn, it is re-introduced in the urn together with a new ball of the same color; thus we will have n colored balls of k different colors and b black balls;
- if a black ball is drawn, a ball of a new color is added to the urn and the new color is assigned a random fitness f_{k+1} drawn from some given distribution:
 - if $f_{k+1} \geq \bar{f}$ (where $\bar{f}$ is the average fitness) a new black ball is also added; thus we will have n colored balls of $k + 1$ different colors and $b + 1$ black balls;
 - if $f_{k+1} < \bar{f}$ no further action is taken; thus we will have n colored balls of $k + 1$ different colors and b black balls;

As to the calculation of $\bar{f}$, we use both a simple average over the colors without weighing them with the number of balls, or a weighted average in which every color is attributed a weight consisting of the number of balls with that color.

We have, therefore, two stochastic processes going on in parallel, one governing the number of black balls: $B_0 = 1$, $B_1 = 1$ or 2 , $B_2 = 1, 2$ or 3 and the other that we have already analyzed, determining the color of the non-black ball added into the urn at time n : $X_1 = 1$, $X_2 = 1$ or 2 , $X_2 = 1, 2$ or 3 , where the probability of adding a new color at time n is given by the number of black balls divided by the total number of balls, i.e. $B_{n-1}/(B_{n-1} + n - 1)$.

In the next sections, we provide numerical results on the behavior of this urn, starting, however, with two limit cases we will use as benchmarks.

3 Two benchmark cases

We first consider two limit cases. First, the case of a standard *Hoppe-Pólya* model, where there is only one black ball acting as mutator. Second, the opposite case in which one new black ball is added whenever a black ball has been drawn, irrespective of any fitness measure. We will refer to this second case as the *constant innovation* model.

In the case of a *Hoppe-Pólya model* innovation tends to disappear quickly, because the probability of extracting a black ball decreases in time as $\frac{1}{t}$. This is exemplified by looking at the time series of the number of colors in one simulation, as reported at the top panels of Fig. 3. A time series of the number of colors is an aggregate measure of the urn variety. At any time step, we can look into the composition of the urn, by considering how the balls of different colors are distributed. In the bottom left panel of Fig. 3, we report the number of balls for each color (mass vector). Here and in the following graphs, colors are numbered by their order of arrival (in this case, the color that is by far mostly represented in the urn is the one that was introduced second). Finally, in the bottom-right panel, we report the frequency histogram corresponding to the allelic partition. In this example we observe a relatively low variety, with two colors scoring more than one thousand balls, and, in particular, color number two getting more than half the total population.

Figure 3: One simulation of the *Hoppe-Pólya* model. Top-Left: time series of the number of colors. Top-Right: time series of the number of colors in logarithmic scale. Bottom-Left: colors mass vector (each color is labeled with an integer number indicating the order of arrival, while the vertical axis reports the number of balls of each color). Bottom-Right: colors allelic partition in log-log scale (the vertical axis reports the frequency of a given number of balls).

If we run the modified model with *constant innovation*, we obtain a completely different scenario. The results of a typical simulation are reported in Fig. 4. The growth rate of the number of colors is slightly lower than linear, as the time series in the left panel suggests. Black balls (the mutator) follow the same dynamics of any other color in this modified model. We know from Arthur et al. (1987) that in a pure *Polya model* a limit value for the relative fractions of colors exists, although this is not known *a priori*. This is not the case with the modification of the *constant innovation model*: whenever the black ball is selected, two balls are added (a new color and another black), so in this case the increase in the total number of balls is twice as large as in the case in which a non-black ball is drawn. Overall, the total number of balls increases faster than t (but still slower than $2t$). As a consequence, the fraction of each color does not converge to a positive limit value, but decreases steadily. Intuitively, this makes sense, because the system is expanding not only in size, but also in variety, and existing colors have to leave space for new colors. In particular, the share of black balls decreases over time, and therefore the innovation rate decreases in the *constant innovation* model too, although more slowly than in the *Hoppe-Pólya* model.

Let us consider the color mass vector (Fig. 4, bottom-left panel) and the color frequency vector (bottom-right panel) of the urn. If we compare with the case of the pure *Hoppe-Pólya* model, there are several colors with a substantial population, and the variety of the system is much larger. The frequency vector of colors in the bottom-right panel is approximately a straight line in a log-log scale, which indicates a power-law distribution for the allelic partition of the *constant innovation model*.

A more systematic way to compare the two models is obtained with a Monte Carlo approach, by running several simulations of the same model specification. Here we run experiments with 1000 repetitions on a time horizon of 2000 steps. Fig. 5 refers to the pure *Hoppe-Pólya* model, and reports the frequency histograms of the final number of colors (left panel) together with

Figure 4: One simulation of the *constant innovation* model. Top-Left: time series of the number of colors. Top-Right: time series of the number of colors in logarithmic scale. Bottom-Left: colors mass vector (each color is labeled with an integer number on the horizontal axis, while the vertical axis reports the number of balls of each color). Bottom-Right: colors allelic partition in log-log scale (the vertical axis reports the frequency of a given number of balls).

the entropy of the system (right panel).

Figure 5: Batch simulations of the *Hoppe-Pólya* model. Distributions of 1000 runs after 2000 steps. Left: frequency histogram of the number of colors. Right: frequency histogram of the entropy.

Fig. 6 contains the results of an identical simulation experiment for the *constant innovation* model. We must keep in mind that, in this limit case, the black color is no different from other colors, and the stochastic process of the number of black balls obeys exactly the same self-reinforcing dynamics. With respect to the pure *Hoppe-Pólya* model, we have by far many more colors. Moreover, the distributions of the number of colors are slightly skewed in opposite ways in the two cases. The most remarkable difference between the two specifications of the model is in the final entropy values. The model with *constant innovation* presents a very skewed

Figure 6: Batch simulations of the *constant innovation* model. Distributions of 1000 runs after 2000 steps. Left: frequency histogram of the number of colors. Right: frequency histogram of the entropy.

distribution, while the distribution is fairly symmetric in the *Hoppe-Pólya* case. In addition, the two distributions have approximately the same variance, but the tails of the two distributions are quite different. For the *Hoppe-Pólya model* the distribution of entropy values is relatively broad (platykurtic), while it is very peaked (leptokurtic) in the *constant innovation* model. Although the modification of the latter brings a much larger variety, as expected, the dispersion of outcomes is lower than in the *Hoppe-Pólya* model.

4 Adaptive urns

In the model with adaptive urns, a newly drawn color is assigned a fitness value from a uniform random distribution, $f_i \sim U[0, 1]$, and the mutator is “rewarded” with a new black ball only if such a fitness value is larger than the average fitness of balls in the urn, $f_i > \bar{f}$. Two specifications of the *adaptive urn* model are considered: one where we use the simple average of extracted fitness values across colors (i.e. if there exist k colors, we average over k fitness values, without weighting the number of balls of each color), and an alternative one where the average fitness is weighted by the number of balls of each color. We call these two versions of the model *simple-average-fitness* and *weighted-average-fitness*, respectively.

We believe that both alternative specifications are informative. If, for instance, we are dealing with technologies, the first specification, *simple-average-fitness*, refers to the case in which technologies can be attributed a measure of intrinsic technological performance independent from the number of users. In other words, the choice of the alternative to be adopted looks only at the characteristics (e.g. performance) of the technology and there is no network or bandwagon effect. Alternatively, this specification can describe the evolution of the technological space as additions of new technologies (patents) to an already existing patent class or the creation of a new patent class. (Cf. the evolution of the technology codes system of the US Patent Office database studied by Strumsky et al. (2012) and Youn et al. (2014).)

The *weighted-average-fitness* specification, on the other hand, reflects cases where the performance measure of each technology/product is weighted by the number of its adopters. This may reflect the existence of some form of network or bandwagon effects but also cases in which products are weighted for instance by the size of their international trade. (Cf. Hausmann and Hidalgo (2011) for the general case or Muro et al. (2011) for the case of *green* technologies.)

Let us begin with the *simple-average-fitness* model. This model is appropriate if one thinks of balls as technologies or products or species. Fig. 7 reports the results of one simulation run lasting for $T = 2000$ steps. First of all, the number of colors increases faster than in the

Figure 7: One simulation of the *Simple-average-fitness* model. Top-Left: time series of the number of colors. Top-Right: time series of the number of colors in logarithmic scale. Middle-Left: colors mass vector (each color is labeled with an integer number on the horizontal axis, while the vertical axis reports the number of balls of each color). Middle-Right: colors allelic partition in log-log scale (the vertical axis reports the frequency of a given number of balls). Bottom-Left: input fitness distribution (Uniform). Bottom-Right: output fitness distribution

Hoppe-Pólya model, but slower than in the *constant innovation* model (top panels). Moreover, we notice that periods without innovations are followed by periods of relatively frequent innovations. The explanation of this behavior is as follows. Innovations alter the average fitness, which is the reference point for an innovation to be considered as ‘successful’. A successful innovation raises this reference point, while unsuccessful innovations do the opposite. On top of this, the main factor driving the dynamics of the average fitness is the self-reinforcing process of ball extraction, which is proportional to the frequency of colors. If a color with relatively low fitness happens to become dominant, the probability of successful innovation increases. If instead a color with relatively high fitness becomes dominant, innovation becomes more and more difficult. In the former case, successive waves of innovation may occur, while in the latter case the dynamics present long periods without innovations.

The middle panels of Fig. 7 contain the color vector (left) and the allelic partition (right) at time $T = 2000$. We observe a richer scenario than in the pure *Hoppe-Pólya* model, with a number of different colors being characterized by more than one ball, but still much less variety

than in the *constant innovation* model. The frequency distribution of the number of balls of each color (middle-right panel), the allelic partition of the urn, is still linear in log-log scale, but it presents a smaller value of the intercept with the horizontal axis with respect to the *constant innovation* model.³

The bottom panels of Fig. 7 report the ‘input’, or initial, fitness distribution (bottom-left panel), and the ‘output’, or final, fitness distribution of balls. The input values are a realization of the input distribution of new balls. The output is how balls’ fitness values are distributed at time $t = 2000$. This distribution is not uniform anymore: the model ‘selects’ a small numbers of fitness values as a result of the Polya self-reinforcing process.

In Fig. 8 we report the results for one simulation run of the model in the *weighted-average-fitness* specification. If we are to compare the results of simulations in Figures 7 and 8, their

Figure 8: One simulation of the *weighted-average-fitness* model. Top-Left: time series of the number of colors Top-Right: time series of the number of colors in logarithmic scale. Middle-Left: colors mass vector (each color is labeled with an integer number on the horizontal axis, while the vertical axis reports the number of balls of each color). Middle-Right: colors allelic partition in log-log scale (the vertical axis reports the frequency of a given number of balls). Bottom-Left: input fitness distribution (Uniform). Bottom-Right: output fitness distribution

³Notice that simulations with the same time horizon and the same number of innovations may have a different final number of balls, because successful innovation events add two balls instead of one.

qualitative behavior is quite similar. The former produces a larger number of colors, as shown by the time series in the top panels, and with the allelic partition in the middle-right panels. Such slightly different outcome is not due to the different specifications of the model but to the path-dependent dynamics of the process, heavily influenced by early events of black ball extraction. We have noticed that simulations of the *Simple-average-fitness* model are more stable, while the *weighted-average-fitness* model presents a larger variability in terms of the final number of balls and their distribution of fitness values. This is due to the sharper measure of average fitness of the *Simple-average-fitness* model, as we will see next with batch simulation experiments.

The time series of the number of colors in the urn as reported by the upper panels of Figures 3, 4, 7 and 8 can be compared to empirical data for the number of technology codes in patents of the USPTO database, for instance. As in Youn et al. (2014), a realistic feature of our results is a diminishing rate of arrival of new colors (technology codes) beside a constant production of balls (patents).

The model with adaptive urns just described may be able to match the right ‘rate’ of technological change intended as the growth rate of colors, since neither the pure *Hoppe-Polya* model or the *constant innovation* model seem to provide a plausible rate. The *adaptive urn* model, lying between these two limit cases, can capture a number of real instances of stochastic growth processes. The *Hoppe-Polya* model expresses a ‘rich-get-richer’ mechanism of self-reinforcement, and in particular can be seen as a *Bose-Einstein* condensate where only one state captures a macroscopic portion of the population (Bianconi and Barabasi, 2001). In the *constant innovation* model, such a mechanism is very much diluted by innovation, which brings an ever increasing segmentation and diversity of the population. The *adaptive urn* offers an intermediate (or metastable) scenario, where the population does not ‘condensate’ in an attracting variety, but nonetheless a structure emerges in the form of a partially segmented population.

Two aspects must be stressed here about path dependence in the model. The first one has to do with the self-reinforcing dynamics of Markov chains driving the competition of alternative options, which obeys the probabilistic rules studied extensively by Arthur et al. (1987). In their model, as in ours, there is a limit value of the relative frequencies, but this limit is not known *a priori*. What makes the difference in our model is that competing options are not given and fixed, but increase in number due to innovation. Also innovation is path-dependent, as in the constant innovation benchmark case, and its pace depends on its own realizations. But in our model, innovation is also “adaptive”, since it depends on the average fitness of alternatives. This dependence has the connotation of a negative feedback: whenever high fitness variants are created, the number of black balls tends to increase fast. But this ‘prepares’ the ground for low innovation rates in following steps, because the average fitness goes up and successful innovations become less likely.

We now turn to batch simulation experiments for both specifications of the adaptive urns model. Beside the *number-of-colors*, we now also look at the *simple-average-fitness* and the *weighted-average-fitness* of the urn. These measures represent the reference value for a ‘successful’ innovation in the two different specifications of the model. Let us start with the *simple-average-fitness* model. Fig. 9 reports the distribution of simulation outcomes at the end of a time horizon of 2000 time steps.

The results are remarkably different with respect to the model without fitness values. The distribution of the number of colors (left panel) is highly skewed: most simulations end with a small number of colors. These results should be compared to Fig. 5 (*Hoppe-Polya* model) and Fig. 6 (*constant innovation* model). As expected, a selection mechanism based on fitness generates a lower number of colors with respect to the *constant innovation* model. If the first colors introduced in the early stages of the simulation are ‘unsuccessful’, the fraction of black balls goes down relatively rapidly, the innovation rate is small and so is the generated variety.⁴

⁴We have run batch simulations of the same model with different time horizons, and found almost

Figure 9: Batch simulations (1000 runs) of the *simple-average-fitness* model with time horizon $T = 2000$ steps. Left: frequency histogram of the number of colors. Center: frequency histogram of the average fitness of colors. Right: frequency histogram of the weighted average fitness of colors.

The simple average fitness distribution is very concentrated around the mean fitness value $1/2$ (Fig. 9, central panel), while the distribution of the weighted average fitness is much broader (right panel). Both features can be empirically relevant, with the *simple-average-fitness* model pointing to systems with relatively low variety, and the *weighted-average-fitness* being more appropriate for scenarios with higher variety.

The batch simulations results for the *weighted-average-fitness* model are reported in Fig. 10. The final distributions of simple and weighted average fitness in the two models are very

Figure 10: Batch simulations (1000 runs) of the *weighted-average-fitness* model with time horizon $T = 2000$ steps. Left: frequency histogram of the number of colors. Center: frequency histogram of the average fitness of colors. Right: frequency histogram of the weighted average fitness of colors.

similar. What is quite different is the distribution of the final number of colors. The *simple-average-fitness* presents a lower dispersion of outcomes, expressed by a lower mean, a lower standard deviation and also a less skewed distribution (Fig. 9, left panel), if compared to the *weighted-average-fitness* model. The reason is that, since the simple average fitness measure is less dispersed, it tends to be a “sharper” reference point for the definition of ‘successful’ innovations (i.e. innovations the fitness of which is precisely greater than this average). Thus, we can expect a lower dispersion of innovation rates across different simulation runs. This translates into less variety for the *simple-average-fitness* model, as shown by the distribution of the final number of colors: for instance, the highest number of colors obtained with the simulation experiment reported in Fig. 9 (*simple-average-fitness* model) is 180, against 250 colors in the experiment of Fig. 10 (*weighted-average-fitness*).

A suggestive interpretation of our results lies in a statistical description of the growth process of firms (Penrose, 1958; Sutton, 1997). Our model can describe the evolutionary dynamics of a given industrial sector, where different colors refer to single firms, and the number of balls of

identical distributions of simple and weighted average fitness, which indicates a stable probability distribution for the stochastic process of the model.

each color is a measure of the firm size (number of employees, or sales volume). At a smaller scale, the model can address the growth process of single firms, as in Bottazzi and Secchi (2006). Within this interpretation, different colors are different submarkets of a firm, and the number of balls of each color measures the size of that particular submarket. Different simulation runs correspond to different firms, and the collection of those simulations in a histogram can be seen as a description of an industrial sector (left panels of Fig. 9 and Fig. 10).

5 Non-uniform fitness distribution

In this section we study the *adaptive urn* model with a non-uniform distribution of fitness values. We use a Beta distribution, for which the probability density function reads as follows:

$$\Phi(f; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} f^{\alpha-1} (1-f)^{\beta-1}, \quad f \in [0, 1]. \quad (3)$$

The factor $\frac{1}{B(\alpha, \beta)}$ is a constant, defined by $B(\alpha, \beta) = \int_0^1 t^{\alpha-1} (1-t)^{\beta-1} dt$. The parameters α and β control the probability distribution, the density function of which can be increasing, decreasing or non-monotonic. Accordingly, the cumulative distribution function $\Phi(f)$ can be convex, concave or S-shaped. The uniform distribution is a special case obtained with $(\alpha = 1, \beta = 1)$. We consider the following cases: $(\alpha = 1, \beta = 3)$, $(\alpha = 2, \beta = 2)$, $(\alpha = 3, \beta = 2)$, $(\alpha = 3, \beta = 1)$ (Figure 11). The mean values are given by $\alpha/(\alpha + \beta)$, and are $1/4$, $1/2$, $3/5$ and $3/4$, respectively.

Figure 11: Beta distribution density function, four cases.

We study the model with a Monte Carlo approach, as before. In each condition, we run 1000 simulation runs for 2000 periods, and look at the final distributions of *number of colors*, *simple average fitness* and *weighted average fitness*. In the case of the model that uses the simple average fitness as a reference value for a successful innovation, the *simple-average-fitness* model, we obtain the results reported in Figure 12. We observe the following. First, whenever the initial distribution of fitness values has more probability mass on larger values, we obtain, on average, a larger number of colors in the urn at any given time, with a larger dispersion of outcomes (Fig. 12, top panels). The distribution remains positively skewed, though. Second, the probability mass of the initial fitness values drives the location of the final distribution both of simple and of weighted average fitness values. For example, with an initial probability distribution $Beta(1, 3)$, we have a mean average fitness equal to 0.25 ± 0.04 , with $Beta(2, 2)$ we get 0.50 ± 0.05 , with $Beta(3, 2)$ we get 0.60 ± 0.04 while with $Beta(3, 1)$ we obtain a mean at 0.75 ± 0.04 . The same mean values are obtained in terms of weighted average fitness, but with much larger variance (almost thrice as much). It is important to notice how the mean values are exactly the theoretical mean values of the starting fitness distribution.

A third observation is about the shape of the weighted average fitness distribution. While a uniform initial distribution of fitness values gives a nearly Gaussian distribution, non-uniform distributions make the final outcomes deviate both in terms of skewness and kurtosis. Except

Figure 12: *Simple-average-fitness* model with non-uniform distribution of fitness values: Left: $Beta(1,3)$. Center-Left: $Beta(2,2)$. Center-Right: $Beta(3,2)$. Right: $Beta(3,1)$. Histograms report the *number of colors*, *simple average fitness* and *weighted average fitness* after $t = 2000$ periods for 1000 simulation runs.

for the $Beta(2,2)$ initial distribution, which is symmetric, the other three cases produce a leptokurtic distribution of final weighted average fitness values. Moreover, a non-symmetric probability mass makes the final distribution skewed, either positively, or negatively, depending on whether the mode of the starting distribution is below or above the mean.

In Figure 13 we report the batch simulations results for the *weighted-average-fitness* model. By comparing these results with the results for the *simple-average-fitness* model in Figure 12, we see some differences in the distributions of the final number of colors. The *weighted-average-fitness* model produces a larger number of colors on average (top panels in Fig. 13) for all four initial Beta distributions. Similarly, the variance is larger. This leads to distributions of the final number of colors in the urn that are more skewed in the *weighted-average-fitness* model than in the *simple-average-fitness* model, a feature that we observe with a uniform distribution of starting fitness values (Figures 9 and 10). Regarding the distributions of simple and weighted average fitness, the two models give very similar results.

6 Fitness-based selection of colors

In this section, we present the final and most complete extension of our urn scheme model with innovation, where not only the probability of the introduction of new colors but also the probability of enforcement of the existing colors is dependent upon fitness values. Relative fitness drives both the novelty generation process (reinforcement of the black ball) and the selection process (reinforcement of the colored balls) and the resulting dynamics is quite rich and interesting.

In the previous versions of the model any color had a probability of extraction equal to the fraction of balls of that color in the urn, while its fitness was only used to reinforce the innovation mechanism (the black ball). In this section we use fitness values also to modify the

Figure 13: *Weighted-average-fitness* model with non-uniform distribution of fitness values: Left: $Beta(1,3)$. Center-Left: $Beta(2,2)$. Center-Right: $Beta(3,2)$. Right: $Beta(3,1)$. Histograms report the *number of colors*, *simple average fitness* and *weighted average fitness* after $t = 2000$ periods for 1000 simulation runs.

color extraction mechanism. For each color i , already present in the urn, we define its extraction probability as follows (Bianconi and Barabasi, 2001):

$$\pi_t^i = \frac{x_t^i f^i}{\sum_{j=1}^{n_t} x_t^j f^j}, \quad (4)$$

where x_t is the color fraction at time t , f^i its fitness, and n_t the number of colors at time t . Such probability is an average fitness measure weighted on the relative abundance of colors, and expresses the fact that both fitness and frequency together determine the extraction probability of a color.

Let us first compare the *simple-average-fitness* and *weighted-average-fitness* specifications of the model equipped with the new probability of color extraction as given by Eq. (4). Figure 14 reports batch simulations which, apart from this new probability of ball addition, use the same settings of previous sections. In both model specifications we obtain a much larger number of colors if compared to the models discussed in the previous section.

The distribution of the final number of colors is not decreasing anymore as it was in Figures 9 and 10 and presents a peak at a relatively large number of colors, between 600 and 700. In the *simple-average-fitness* model (left panel of Fig. 14) the distribution is unimodal, with about 660 colors being the most frequent outcome. This is not the case for the *weighted-average-fitness* model, which gives a bi-modal distribution. In this case, a relatively large number of simulations end up with few colors, and the second peak is less pronounced. The mode is also slightly smaller, with a number of colors around 650 being the most frequent one. In the *weighted-average-fitness* model there is a threshold effect in the color generation mechanism of the adaptive urn process. Innovation (the black ball) is rewarded if the fitness of new colors is above the weighted average. But if such new colors become too abundant, the weighted average increases, the innovation rate declines and the number of colors increases more slowly.

Figure 14: Fitness-based selection of colors as from Eq. (4). Left: *simple-average-fitness* model. Right: *weighted-average-fitness* model. Histograms collect 1000 simulation runs, each lasting 2000 periods.

This effect is absent in the *simple-average-fitness* model, where the relative weight of colors plays no role in the innovation reinforcement mechanism. In this case, the fitness dependent selection mechanism (equation 4) boosts the diversity of the system, while in the *weighted-average-fitness* model this occurs only partially, because the increasing share of high fitness balls makes successful innovation less and less likely.

Finally, we introduce a parameter γ that controls for the relative strength of fraction and fitness factors, and redefine the colors extraction probability as

$$\pi_t^i = \frac{(x_t^i)^\gamma (f^i)^{(1-\gamma)}}{\sum_{j=1}^{n_t} (x_t^j)^\gamma (f^j)^{(1-\gamma)}}. \quad (5)$$

This definition has two useful features. First, it includes the previous versions of the model of Sections 4 and 5, which is obtained with setting $\gamma = 1$ (while with $\gamma = 0$ the extraction probability is proportional to the color fitness). Second, the new parameter could allow us, in principle, to calibrate the model on empirical data.

We analyze this last specification of the fitness-based *adaptive urn* model by focusing on the *simple-average-fitness* version. Figure 15 reports batch simulation results in terms of the distribution of the final number of colors. When γ is low, the fitness value prevails in the colors

Figure 15: Fitness-based selection of colors given by Eq. 5 (only *simple-average-fitness* model). γ weights the contribution of fitness to the colors extraction probability. Left: $\gamma = 0.25$. Center: $\gamma = 0.5$. Right: $\gamma = 0.75$. Histograms collect 1000 simulation runs. Each run lasts for $T = 2000$ periods.

extraction probability (Eq. 5). This condition produces a relatively larger number of colors on average, with relatively less dispersion (top-left panel of Figure 15). A low γ , on the contrary,

puts more weight on the colors fraction than on their fitness, which results in a lower number of colors on average, with more dispersion (bottom-left panel of Figure 15). Increasing γ , we take the model closer to the specification of the previous sections, where the colors extraction probability was only determined by their relative share. In terms of entropy, we notice no differences for what concerns the mode of the distribution between $\gamma = 0.25$, $\gamma = 0.5$ and $\gamma = 0.75$, which is slightly above 7 in all three cases. A larger γ increases only marginally the dispersion of results, with few simulations ending with a substantially lower entropy.

To summarize, the specification of Eq. (5) enables one to tune the model by giving more or less importance to the fraction or the fitness of colors by changing the parameter γ .

Finally, figure 16 shows the color mass distribution of typical runs of the fitness-based selection model when fitness values are drawn, respectively, from a uniform distribution, a Beta distribution with $\alpha = 1$ and $\beta = 3$, and a Beta distribution with $\alpha = 3$ and $\beta = 1$. The graphs show that, although path-dependence remains high, the population is not entirely dominated by colors introduced in the very first draw, and also that colors arriving between time 10 and 20 gain a relevant share in the population.

Figure 16: Color mass vector of a typical single run for fitness-based selection. Left: uniform fitness distribution $\gamma = 0.5$. Center: beta fitness distribution $\alpha = 1$ and $\beta = 3$. Right: beta fitness distribution $\alpha = 3$ and $\beta = 1$.

7 Conclusions

The phenomenon of path dependence has attracted the attention of many scholars interested in the evolution of technologies (David, 1985; Arthur, 1989) or of social, economic and political institutions (David, 1985; Pierson, 2000). However, the conclusions supporting path dependence and lock-in and coming both from empirical evidence and from analytical models have also been questioned and criticized (Liebowitz and Margolis, 1995). The criticisms of Liebowitz and Margolis are raised from a neoclassical perspective on the grounds of the belief that inefficient allocations cannot be persistent and will be always remedied by the market. In this paper, we have, by way of contrast, attempted to provide some tools that could extend and somehow qualify the standard models of path dependence on the grounds of a more comprehensive evolutionary perspective.

Standard models based on Pólya urns only consider a selection process among given alternatives and do not allow for the arrival of novelty and variation. In this paper, we introduced some extensions of urn models in which, in addition to the stochastic process governing selection, there exists another stochastic process for novelty generation. Paraphrasing Liebowitz and Margolis, we could say that lock-in into a supposedly inferior technology such as VHS (as opposed to Betamax) is ultimately itself going to be transient not so much because market efficiency will take its revenge but because both VHS and Betamax may finally be displaced by the arrival of the DVD technology. The models we simulated in this paper can account for this complex interplay between selection among the alternatives available in some given time span and the arrival of new alternatives, which changes the domain on which selection operates.

In this paper, we have analyzed by means of numerical simulations the properties of different specifications of urn models with stochastic additions of new colors. We believe that these

models can be calibrated in such a way as to give us useful insight on real cases of evolution of technologies, organizations and institutions whereby phases in which path dependent selection among given alternative are followed by phases in which the dynamics is mainly driven by the introduction of new alternatives.

Here we focused on the methodology of our urns with adaptive mutation rate and we leave detailed empirical applications to future work. However, already in this abstract formulation the model can give useful insight on several real processes presenting the co-occurrence of innovation, adaptive selection and self-reinforcing mechanisms. For instance, our model can describe birth and growth processes subject to innovation and positive feedback such as cities or firms. In this interpretation, fitness can represent a size measure. In order to build on this interpretation of the model, a statistical characterization of the number of colors distribution is needed. If one looks at the size distribution of cities, the models parameters should be set in order to obtain the empirical Zipf law distribution (Gabaix, 1999). Alternatively, the *adaptive urn* model can be used to explain the statistical features of firm growth rates (Bottazzi and Secchi, 2006). In any case, the initial distribution of fitness values is one factor where parameters can be set in order to match the empirical distribution of interest with the model. A second factor is the choice between the weighted and the simple average fitness specification of the model for the reference value of successful innovation: if more variability is needed, the *weighted-average-fitness* model is more appropriate, while for relatively lower variability scenarios the best choice is the *simple-average-fitness* model.

To conclude, our model has explored the potential of stochastic growth processes with an adaptive interplay between selection and innovation. We believe this probabilistic framework stands as a powerful model of real world dynamic systems that show neither the homogeneity of an attracting state or a fully heterogeneous picture, but instead present the emergence of intermediate metastable structures.

References

- ALEXANDER, J. M., B. SKYRMS, AND S. ZABELL (2012): “Inventing new signals,” *Dynamic Games and Applications*, 2, 129–145.
- ARTHUR, B. (1989): “Competing technologies, increasing returns, and lock-in by historical events,” *Economic Journal*, 99, 116–131.
- ARTHUR, W., Y. ERMOLIEV, AND Y. KANIOVSKI (1987): “Path-dependent processes and the emergence of macrostructure,” *European Journal of Operation Research*, 30, 294–303.
- BIANCONI, G. AND A.-L. BARABASI (2001): “Bose-Einstein Condensation in Complex Networks,” *Physical Review Letters*, 86, 5632–5635.
- BOTTAZZI, G. AND A. SECCHI (2006): “Gibrat’s law and diversification,” *Industrial and Corporate Change*, 15, 847–875.
- CRAWFORD, V. AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1451.
- DAVID, P. (1985): “Clio and the economics of QWERTY,” *American Economic Review*, 75, 332–337.
- DOSI, G., Y. ERMOLIEV, AND Y. KANIOVSKI (1994): “Generalized urn schemes and technological dynamics,” *Journal of Mathematical Economics*, 23, 1–19.

- EREV, I. AND A. ROTH (1998): “Predicting how people play games: reinforcement learning in experimental games with unique, mixed-strategy equilibria,” *American Economic Review*, 88, 848–881.
- EWENS, W. J. (1972): “The sampling theory of selectively neutral alleles,” *Theoretical Population Biology*, 3, 87–112.
- GABAIX, X. (1999): “Zipf’s law and the growth of cities,” *American Economic Review*, 89, 129–132.
- HAUSMANN, R. AND C. A. HIDALGO (2011): “The network structure of economic output,” *Journal of Economic Growth*, 16, 309–342.
- HOPPE, F. (1984): “Polya-like urns and the Ewens’ sampling formula,” *Journal of Mathematical Biology*, 20, 91–94.
- KAUFFMAN, S. A. (1993): *The Origins of Order*, Oxford: Oxford University Press.
- LIEBOWITZ, S. AND S. MARGOLIS (1995): “Path Dependence, Lock-in and History,” *Journal of Law, Economics and Organization*, 11, 205–226.
- MURO, M., J. ROTHWELL, AND D. SAHA (2011): “Sizing the Green Economy: a national and regional green jobs assessment,” Working paper, metropolitan policy program, Brookings Institution.
- PENROSE, E. (1958): *The Theory of the Growth of the Firm*, Oxford: Oxford University Press.
- PIERSON, P. (2000): “Increasing Returns, Path Dependence, and the Study of Politics,” *American Political Science Review*, 92, 251–267.
- STRUMSKY, D., J. LOBO, AND S. VAN DER LEEUW (2012): “Using patents technology codes to study technological change,” *Economics of Innovation and New Technology*, 21, 267–286.
- SUTTON, J. (1997): “Gibrat’s legacy,” *Journal of Economic Literature*, 35, 40–59.
- TRIA, F., V. LORETO, V. SERVEDIO, AND S. STROGATZ (2014): “The dynamics of correlated novelties,” *Scientific Reports*, 4, 05890.
- YOUN, H., L. BETTENCOURT, D. STRUMSKY, AND J. LOBO (2014): “Invention as a Combinatorial Process: Evidence from U.S. Patents,” Working paper, Institute for New Economic Thinking, Oxford University.
- ZABELL, S. L. (1992): “Predicting the unpredictable,” *Synthese*, 90, 205–232.