

Staccioli, Jacopo; Virgillito, Maria Enrica

Working Paper

Back to the past: The historical roots of labour-saving automation

LEM Working Paper Series, No. 2020/34

Provided in Cooperation with:

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

Suggested Citation: Staccioli, Jacopo; Virgillito, Maria Enrica (2020) : Back to the past: The historical roots of labour-saving automation, LEM Working Paper Series, No. 2020/34, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/243492>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

INSTITUTE
OF ECONOMICS



Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

**Back to the past: the historical roots of
labour-saving automation**

Jacopo Staccioli §¶
Maria Enrica Virgillito ¶§

§ Department of Economic Policy, Università Cattolica del Sacro Cuore, Milan, Italy
¶ Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy

2020/34

November 2020

ISSN(ONLINE) 2284-0400

Back to the past: the historical roots of labour-saving automation *

Jacopo Staccioli^{†‡§}

Maria Enrica Virgillito^{††}

21st November 2020

Abstract

This paper, relying on a still relatively unexplored long-term dataset on U.S. patenting activity, provides empirical evidence on the history of labour-saving innovations back to early 19th century. The identification of mechanisation/automation heuristics, retrieved via textual content analysis on current robotic technologies by Montobbio et al. (2020), allows to focus on a limited set of CPC codes where mechanisation and automation technologies are more prevalent. We track their time evolution, clustering, eventual emergence of wavy behaviour, and their comovements with long-term GDP growth. Our results challenge both the general-purpose technology approach and the strict 50-year Kondratiev cycle, while provide evidence of the emergence of erratic constellations of heterogeneous technological artefacts, in line with the development-block approach enabled by autocatalytic systems.

JEL classification: O3, C38, J24.

Keywords: Labour-Saving Technologies, Search Heuristics, Industrial Revolutions, Wavelet analysis.

*Maria Enrica Virgillito acknowledges support from European Union's Horizon 2020 research and innovation programme under grant agreement No. 822781 GROWINPRO – Growth Welfare Innovation Productivity.

[†]Department of Economic Policy, Università Cattolica del Sacro Cuore, Via Necchi 5 – 20123 Milano, Italy.

[‡]Institute of Economics and EMbeDS, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33 – 56127 Pisa, Italy.

[§]To whom correspondence should be addressed: ✉ jacopo.staccioli@unicatt.it

1 Introduction

The existence of labour-saving (hereafter, LS) heuristics driving the rate and direction of technological change is a documented pattern, since the inception of the First Industrial Revolution. Reducing the time of operations during Taylorism, increasing the saturation of *takt*-times during Toyotism, and speeding up processes and executions of functions remotely tracking operators' intervention nowadays, are the core drivers of mechanisation and automation.

In the tradition of the economics of innovation, the First Industrial Revolution had been a combination of *time-saving heuristics*, enabled by the mechanisation process, and the division of labour inside factories, together with the emergence of innovative artefacts. The role played by time-saving heuristics in shaping the direction of mechanisation has been emphasised by von Tunzelmann (1995) with reference to the cotton industry in the British Industrial Revolution: the massive increase in labour productivity resulted from the use of innovation and discovery through which a spinner was able to produce in a day as much yarn as previously required by a full year of work, without mechanisation.¹

On top of that, Freeman (2019) conceptualises the First Industrial Revolution as a paradigmatic shift emerging from the combination of time-saving heuristics on the one hand, and the new clear demarcation between working- and life-time for wage labourers on the other hand, an attitude absent in the pre-industrial societies (Thompson, 1963), allowing workers discipline and ensuring their participation to the productive activities, e.g. by turning Monday into a working, rather than a drinking, day. As corroborating evidence, using a detailed and quite granular report, the Hand and Machine Labor Study commissioned by the Department of Labor in 1899 to detect the impact of mechanisation on labour productivity, Atack et al. (2020) estimate that only one-third of the increase in labour productivity (measured as time spent in a given operation) in the late nineteenth century was due to 'inanimate power', while the rest unexplained component remains attributed to other factors, among which division of labour plays a prominent role.

Speeding up the production process clearly maps into the need of reducing the human active participation to the process itself. Therefore, time-saving and LS heuristics have been considered by economic historians as potential *focussing devices* (Rosenberg, 1976) guiding the search process, however of a very particular type. In general, in the development of a new artefact, inventors face technical trade-offs and bottlenecks which have to be overridden. Search heuristics might have various nature and directions (Cohen et al., 1996), going from the 'make it smaller' for microprocessors, to the 'make it faster' for aircraft, even to the 'make it more exclusive' for smartphones (e.g. Apple's iPhone). Indeed, focussing devices are rather heterogeneous among inventors, and as such they *locally* guide the search and discovery process, defining the technological trajectory, inside a given established paradigm (Dosi, 1982). This is not the case for LS heuristics which, rather than local, appear as a generalised feature of the history of innovation and in general of capitalism.

¹The author refers in particular to Baines (1835).

“In England, strikes have regularly given rise to the invention and application of new machines. Machines were, it may be said, the weapon employed by the capitalists to equal the result of specialised labour. The *self-acting mule*, the greatest invention of modern industry put out of action the spinners who were in revolt. If combinations and strikes had no other effect than of making the efforts of mechanical genius react against them, they would still exercise an immense influence on the development of the industry.”

[Marx (1955, p. 161); also cited in Rosenberg (1976, p. 118)]

Are these LS heuristics empirically detectable? Attempts to infer heuristics and knowledge bases appear e.g. in Castaldi et al. (2009) at the artefact level, focussing on the tank technology and the evolution of its attributes over time, but also in Martinelli (2012), who uses patent-citation networks to infer the emergence of new paradigms by changes in bottlenecks and search heuristics, therefore at the so-called knowledge level. Recently, Taalbi (2017), relying on specialistic trade journals, collected information about drivers of innovative activities supposed to be relevant by innovators, and investigates eventual distinct patterns across industry and over time.

Currently, heuristics are usually inferred by the technical engineering literature and related case-studies. However, patents and their textual content also provide a good source of information to detect codified knowledge and the ensuing search heuristics. Relatedly, the use of textual analysis techniques enables a comprehensive study of large scale textual dataset. By looking at the textual contents of robotic patents over the last decade, Montobbio et al. (2020) are able to isolate those ones which clearly embed a LS trait. The identification of LS patents, done by natural language processing which includes probabilistic topic modelling, lead to a clear definition of the set of technological artefacts behind LS robotic patents published by the USPTO between 2009 and 2018. Two insightful excerpts from LS patents follow:

“Automated systems, such as robotic systems, are used in a variety of industries to **reduce labo[ur] costs and/or increase productivity**. Additionally, the use of human operators can involve increased cost relative to automated systems.” [US20170178485A1]

“The use of the technology [robots] results in improved management of information, services, and data, increased efficiency, significant reduction of time, **decreased manpower requirements**, and substantial cost savings.” [US20100223134A1]

After identifying patents explicitly containing LS heuristics, Montobbio et al. (2020) infer the type of human activities that the technology contained in LS patents is intended to replace, by capturing both the formal technological content of the invention using patent classification codes and the substantial purpose of broader robotic innovations, using the vector of words which characterises each topic. Thanks to this twofold analysis, they describe those fields and activities that are more exposed to LS innovations. LS patents

appear to be concentrated in particular in the following fields: (i) Transport, Storage and Packaging, (ii) Diagnosis and Therapy, (iii) Transmission of Digital Information, (iv) Optical elements, (v) Chemical or Physical Laboratory Apparatus (measuring and testing in chemistry), and (vi) Moving Parts.

The authors propose a taxonomy wherein it emerges that the typical tasks on which LS research effort is focussed include (i) dexterity and manipulations, as in packing, storing, conveying, and handling packages in the logistics industry; (ii) activities entailing social intelligence, such as caretaking patients and the elders; (iii) activities requiring cognitive intelligence and complex reasoning, e.g. the ability of predicting, learning, classifying and evaluating, typical of high-level professional segments. Notably, the analysis shows that the overall bundle of technologies behind LS heuristics is not simply related to robots *stricto sensu*, but it encompasses a wider set of technologies, functions and operations. In this respect, rather than interpreting the new wave of LS technologies as the next GPT (Trajtenberg, 2019), to genuinely account for the unfolding of the latest wave LS technologies, a ‘technological constellation’ perspective à la Freeman would be more informative.

In this paper we intend to move ahead by delving into the past, i.e. by adopting a ‘historical technological constellation perspective’ and looking at the emergence and evolution of the bundle of technologies behind the current LS heuristics detected in robotic innovations. Indeed, as we shall show, mechanisation and automation are not the result of a single dominant product design, but rather of a bundle of technological artefacts, which experience patterns of comovements, anti-comovements, explosion, and dissipation. Our empirical investigation, which looks at historical patent data over the period 1836–2019 vindicates, first, the underlying technological complexity, in terms of bundles of output, behind LS technologies; second, the increasing historical relevance of those technological artefacts entailing mechanisation and automation; last, the absence of a neat recurrence of periodic waves of innovations. In fact, although we identify the emergence of long waves characterising part of these technological artefacts, they are hardly periodically recurrent.

Our findings clearly parallel the Schumpeterian reading of capitalist systems (Schumpeter, 1939) which epitomises the perspective of long waves of technological innovations, with phases of upswings and downswings, clustering of heterogeneous innovations and patterns of interdependence among them, giving rise to upsurge and transformation, alternated with phases of slackening and declines.

Building upon the Schumpeterian perspective, according to Freeman and Louçã (2001) the history of modernisation is punctuated by distinct phases, characterised in terms of dominant technological systems, or better *techno-economic paradigms*. They go beyond the deterministic Kondratiev wave approach put forward by Schumpeter, and propose the notion of ‘constellations of major technical innovations’, by far more complex than the popularised GPT version, entailing the diffusion of unique technologies, say steam, electricity, ICTs, and now AI, which for their pervasiveness encompass all sectors of the economy, and therefore turn out to entirely characterise the process of economic growth.

Differently, constellations pertain to the notion of autocatalytic mechanisms, entailing *development blocks* of technological artefacts (Dahmén, 1988). Therefore, according to this

perspective (Nuvolari, 2019), the development block underlying the British Industrial Revolution consisted of machinery, machine tools, steam engines, coal, and iron production techniques, while the one underlying the Third Industrial Revolution consists of semiconductors, computers, software, and networking equipment.

The periodic cycle approach proposed by Kondratiev, rephrased by Schumpeter, and then endorsed by Perez (1983), has been challenged by the empirical literature and questioned particularly by Silverberg (2007), who highlights a series of drawbacks characterising the empirical detection of long cycles. The latter pertain, first, to the non-stationary nature of long-term time series and, related, to the distortion imposed by making the series stationary using whatever detrending technique; second, to the short time horizon characterising the majority of the analyses, with many series lasting exactly fifty years, and therefore over-imposing the Kondratiev wave; third, to the absence of a dataset for *true* innovations; fourth, to arbitrary trimming of the dataset.

These drawbacks have been taken into account by Silverberg and Verspagen (2003), who dismiss the long cycle perspective and opt for a more neutral detection of clustering of innovations by fitting a Poisson model, under the hypothesis of absence of clustering, versus a negative binomial model, allowing for clustered events, and therefore for a variance component. Although innovation clustering is verified, any periodic deterministic cluster hardly emerges:

“Innovations may indeed cluster, but not in any deterministic sense, and their pattern may shed light on a unified mechanism explaining a range of their properties. Aggregate economic activity, simultaneously with certain patterns of structural change, may obey certain laws that dialectically intertwine chance and necessity and produce robust patterns, but ones that do not lend themselves to any very simple forecasting. It is on this note that I hope long waves will long be with us as a field of scientific research.” [Silverberg (2007)]

In the following, we shall proceed by explicitly addressing the major drawbacks pointed out by Silverberg (2007) by going beyond limited spectral analysis and filtering techniques, and resorting to a non-stationary resilient methodology, namely wavelet analysis, which will be performed upon a well defined set of technological innovations, i.e. patents published since the 1790, and therefore also overcoming the shortness of the data structure, and avoiding to super-impose the long cycle identification. As said above, we focus on a particular subset of the overall technological artefacts, namely those who have been recognised to *currently* involve explicit LS heuristics.

Linking the evolutionary literature studying the employment impact of technical change, theoretically discussing different compensation mechanisms balancing labour-saving effects of innovation (Calvino and Virgillito, 2018; Dosi et al., 2019; Freeman and Soete, 1987; Piva and Vivarelli, 2018; Simonetti et al., 2000; Vivarelli, 1995),² the study of

²Many contributions are emerging in the recent years looking at the impact of automation adopting a neo-classical perspective (Acemoglu and Restrepo, 2018, 2019, 2020; Graetz and Michaels, 2018) mostly relying on sectoral and local labour market analysis. The evolutionary tradition distinguishing heterogeneous impacts of embodied vs disembodied technical change upon employment has been explored in Barbieri et al.,

knowledge bases embedded in technology (Dosi, 1988) and the emergence of long waves or alternatively of clusters of innovations (Silverberg and Verspagen, 2003), our contribution departs from the literature in terms of both the novelty of the empirical analysis, by fully exploiting the long-run historical dimension of the USPTO dataset, still relatively unexplored, the use of wavelet analysis to study patent data, and ultimately enriches our understanding of the long run history of the constellations of artefacts behind current LS robotic technologies. Indeed, in the wake of a purported Fourth Industrial Revolution³, and of the over-abused statement ‘this time is different’, the understanding of the evolution of the bundle of technologies behind the current explicit LS robotic ones might allow a thorough and deeper policy action to counteract labour shedding effects.

The paper is organised as follows: in Section 2 we identify the long-term evolution of the constellation of technologies behind current LS innovations, we present their time trend and clustering patterns. Section 3 detects the presence of temporal cycles in the data by means of wavelet analysis and explores the extent to which the intensity of innovative activity is correlated with business cycles and recessions. Finally Section 4 concludes by outlining potential avenues of further research and useful policy implications.

2 Back to the past: labour-saving innovations since 1830s

The first step of our empirical investigation entails the determination of technological classes which are recognised to currently present LS traits, with the aim at delving into the past and tracing a historical account of their evolution. Our main source is Montobbio et al. (2020), which investigates the presence of LS heuristics within a set of 29,789 robotic patent applications published by the USPTO between 2009 and 2018 and quantitatively identify, through a probabilistic topic model of their full-texts, the CPC (Cooperative Patent Classification) codes which bear the most relevance to underlying LS innovations. In a nutshell, the prevalence of a LS trait is identified by means of a metric pointing at those topics which are more prevalent in LS robotic patents vis-à-vis the population of generic robotic patents. Topics are then matched to CPC codes. The metric used by Montobbio et al. (2020), namely a topic relevance distribution, is also reported here in Fig. 1 for convenience (a more detailed technical summary of Montobbio et al., 2020 is also provided in Appendix A).

In the present paper, we build upon their results by studying the long-term evolution in the attribution by patent examiners of CPC codes linked to LS technology. Starting from Table 1 in Montobbio et al. (2020), we restrict the analysis to the set of 3-digit CPC codes which exhibit a weight of at least 10% (5th column) among topics with positive relevance ($\tilde{\Theta}_k^{LS} > 0$) to LS patents (2nd column), as opposed to general robotic patents.⁴ We also include 4 CPC codes excluded from the said table in that they are widely pervasive of both

2018; Pellegrino et al., 2019; Van Roy et al., 2018 mostly adopting firm-level data with a panel structure, to mention a few recent studies. More established notions of process vs product innovations are explored in (Harrison et al., 2014; Lachenmaier and Rottmann, 2011; Van Reenen, 1997).

³For empirical evidence investigating current Industry 4.0 trends in the automotive industries see Cirillo et al. (2020) and Moro et al. (2019)

⁴Codes which belong to CPC ‘raccord’ class Y are left full digit.




Figure 1: Topic relevance metrics Θ_k^{rob} and Θ_k^{LS} for robotic patents (blue, in descending order) and their LS subset (orange). Source: Montobbio et al. (2020).

LS patents and general robotic patents (see Montobbio et al., 2020, footnote 14): these are B25 (“Hand tools; Portable power-driven tools; Manipulators”), G01 (“Measuring; Testing”), G05 (“Controlling; Regulating”), G06 (“Computing; Calculating; Counting”), and Y10S901, which points to the “Robots” former USPC Class 901. We end up with a set of 15 CPC codes, outlined in Table 1 along with their official definition. In order to investigate their historical use within patenting activity, we employ the Master Classification File (MCF) for U.S. Patent Grants.⁵ The dataset reports, for each patent granted by the USPTO since July 1790, the corresponding CPC classification. In fact, at the time of writing, some records are missing, and usable data start in 1836. Our analysis is therefore restricted to the period 1836–2019. Notably, titles, abstracts and full-texts of USPTO patents are available only since the 1970s; therefore, we resort to CPC codes rather than textual analysis in order to elicit long-run information on LS technologies. By looking at Table 1, it is clear that CPC codes characterising current LS robotic patents are quite diverse and range from more obvious CPC codes, such as ‘Machine tools’ (B23) and ‘Hand tools’ (B25), to less so, such as ‘Optics’ (G02) and ‘Medical of veterinary science’ (A61). Therefore, the so-called constellation approach, as opposed to the GPT one, appears more appropriate to characterise long-run automation trends. Hereafter, we will collectively refer to these CPC codes as target CPCs, granted the degree of heterogeneity therein, but nonetheless conveying important information on the functions and activities involved by the underlying artefacts and the related human functions. We now ask the question: how does patent intensity of target CPCs evolve over time?

Fig. 2 shows the evolution in the assignment of target CPC codes by patent examiners within each year in our time window. The vertical axis measures patent intensity in relative terms, i.e. it represents the number of times the underlying CPC has been assigned

⁵Available here: <https://bulkdata.uspto.gov/data/patent/classification/cpc/>

CPC	Definition
A61	Medical of veterinary science; Hygiene
B01	Physical or chemical processes or apparatus in general
B23	Machine tools; Metal-working not otherwise provided for
B25	Hand tools; Portable power-driven tools; Manipulators
B62	Land vehicles for travelling otherwise than on rails
B65	Conveying; Packing, Storing; Handling thin or filamentary material
C12	Biochemistry; Beer; Spirits; Wine; Vinegar; Microbiology; Enzymology; Mutation or genetic engineering
G01	Measuring; Testing
G02	Optics
G05	Controlling; Regulating
G06	Computing; Calculating; Counting
H01	Basic electric elements
H04	Electric communication technique
Y10S901	Robots
Y10T436	Chemistry: analytical and immunological testing

Table 1: Target CPC codes definitions.




Figure 2: Time evolution of relative assignment frequencies for target CPC codes.

divided by the overall number of assigned CPC codes in the same year, as follows:

$$\text{patent intensity of code } CPC \text{ in year } t = \frac{\text{number of } CPC \text{ assignments in year } t}{\text{number of all assignments in year } t}$$

The vertical axis therefore measures the relative importance of target CPCs vis-à-vis the rest of technological classes. Already at a first glance, a sizeable heterogeneity emerges in their time evolution. Indeed, different CPC codes exhibit starkly different evolutions, both in terms of shape and scale. A few codes have witnessed an overall steeply increasing trend (A61, C12, G02, G06, H01, H04, Y10S901, Y10T436), suggesting a century-long ever growing commitment to innovative effort in the underlying fields, while the remaining display a more ‘stationary’ or even decreasing dynamics. Exploding trends characterise technologies related to computing, processing, and testing, namely pivotal functions at the core of the Third Industrial Revolution. A more steady tendency is instead recorded for those technologies related to codes B65 and G05, mostly characterising innovations dealing with storage, packing, conveying, and regulating control processes. Patent intensity of these technologies seem to be rather persistent over time, in a fashion more akin to enabling technologies, processes, and methods, rather than products themselves. A decaying trend is instead visible for codes B23, B25 and B62, mainly characterising the First and Second industrial revolutions, such as innovation in the agricultural sectors (land vehicles) and related to the mechanisation process in the assembly line, providing hand tools, manipulators, and metal-working.

But how relevant are these target CPC codes when compared to the rest of innovative efforts? In order to better grasp the relevance, in term of assignment frequency, of our target codes vis-à-vis the rest of CPC codes, we divide our time window into five sub-periods and look at the most common classification codes therein. Breakpoints are set at 1880, 1920, 1960, and 2000.

Fig. 3 reports the ranking of the top 10 CPC codes by assignment frequency in each sub-period. Between 1836 and 1880 (panel (a)), innovative effort is primarily devoted to advancements in the primary sector, as captured by CPC code A01 (“Agriculture; Forestry; Animal husbandry; Hunting; Trapping; Fishing”), followed by engineering, after code F16 (“Engineering elements and units; [...]”). Between 1881 and 1920 (panel (b)), the picture witnesses a reversal, with engineering now prevailing over agriculture, as code F16 takes over A01 for first position. The rise of engineering also brings to the fore related applications to transportation, with codes B60 (“Vehicles in general”) and B61 (“Railways”) ranking third and sixth, respectively. This trend intensifies in the next period, from 1921 to 1960 (panel (c)), in which agriculture disappears from the top 10 positions, and engineering (F16) is now trailed by applications to the extraction and energy sector (C10, “Petroleum, gas or coke industries; [...]”). Towards the end of this period, in late 1940s and 1950s, the transistor and its most used design, the MOSFET, are invented, boosting electricity and electrical components to third position. The related CPC code, H01 (“Basic electric elements”), becomes the most commonly assigned code ever since, ranking first in the remaining two sub-periods (panels (d) and (e)).

The transistor revolution opens up for advances in other fields, most notably biotechnology, which consistently ranks second from 1961 to date. What differentiates the 1961–2000 period from the 2001–2019 period is a focus on chemistry in the former, with codes C07 (“Organic chemistry”) and C08 (“Organic macromolecular compounds; [...]”) ranking third and fourth, respectively, and on ICT in the latter, witnessed by codes H04 (“Electric communication technique”) and G06 (“Computing; Calculating; Counting”) in 4th and 5th position, respectively.⁶ Moreover, the relative frequency scale (reported in the horizontal axis) appears to double from one period to the next in the last two sub-periods: the most common code accounts for $\approx 4\%$ of all assignments between 1921 and 1960, for $\approx 8\%$ of all assignments between 1961 and 2000, and for $\approx 16\%$ between 2001 and 2019. This suggests an increasing concentration of innovative activity within the underlying fields over time.

Fig. 3 also captures the overlap between our target CPC codes with respect to other commonly assigned codes. It is immediate to notice that our target codes (highlighted in blue, as opposed to general codes pictured in orange) become increasingly widespread over time, with a remarkable thrust after World War II: half of the top 10 codes assigned between 1961 and 2000 (panel (d)) and 7 after the turn of the Millennium (panel (e)) belong to our list of potentially LS codes.

Finally, Fig. 3 highlights (in green) the degree of co-assignment of LS CPC codes with other CPC codes at the single patent level. The CPC co-occurrence is informative of the degree of technological complementarity and underlying technological complexity. Indeed, over time, not only our target CPC codes increase in relevance, but exhibit a considerable degree of complementarity with other technologies (C07, C08), meaning that potential LS heuristics are also nested in conjunction with other non-explicitly LS CPC codes.

As already noted, the (relative) assignment frequency of different CPC codes follows markedly distinct dynamics over time. A second look at Fig. 2 reveals three main evolution patterns. CPC codes B23, B25, B62, and B65, whose (normalised) evolutions are reported superimposed in Fig. 4(a) for the sake of clarity, are all characterised by an ‘early peak’, in the sense that they reach their maximum frequency of assignment early in our time window, no later than World War II, and exhibit a decreasing trend afterwards. Since this cluster includes, among others, inventions such as power driven tools, manipulators, and transport equipment, it is apparent that effort towards automation of machining tasks chronologically precedes the ‘transistor revolution’ and the ensuing ‘IT revolution’.

Similarly, codes A61, C12, G02, G06, H01, and H04, pictured superimposed in Fig. 4(b), all exhibit a monotone and steeply increasing trend. This suggests that the underlying innovative effort, whose target fields range from bioscience to ICT, and which started to gain momentum during the post-World War II economic boom, has yet to plateau.

Finally, codes B01, G01, and G05 form a third cluster, as the relevant assignment frequen-

⁶For the sake of completeness, other recurrent CPC codes across the five periods include: A47 (“Furniture; Domestic articles or appliances; Coffee mills; Spicemills; Suction cleaners in general”), B29 (“Working of plastics; Working of substances in a plastic state in general”), E05 (“Locks; Keys; Window or door fittings; Safes”), Y10T24 (“Buckles, buttons, clasps, etc.”), Y10T74 (“Machine element or mechanism”), Y10T83 (“Cutting”), Y10T137 (“Fluid handling”).






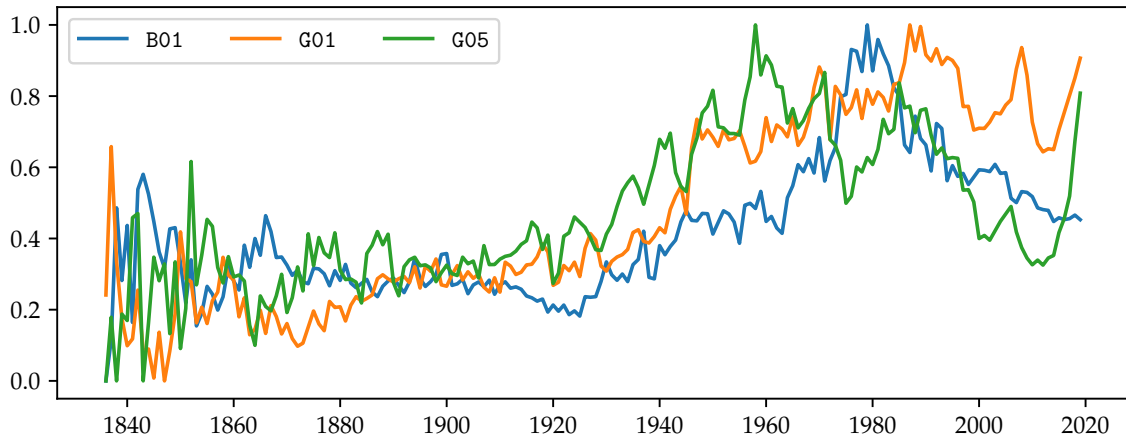
Figure 3: Overall CPC relevance over selected time window, and co-occurrences between target and non-target CPC codes. Blue and orange bars measure the proportion of, respectively, target and non-target CPC codes as a fraction of all CPC codes assigned in the underlying period. Green bars measure the degree of co-occurrence of CPC target codes with specified non-target codes.



(a) 1st cluster: hump-shaped assignment frequency.



(b) 2nd cluster: ever increasing assignment frequency.



(c) 3rd cluster: plateaued assignment frequency.

Figure 4: Three clusters of innovative activity. All series are defined as in Fig. 2 and normalised to span the unit interval.

cies, pictured superimposed in Fig. 4(c), are overall increasing, but underlying technologies seem to plateau in the second half of the 20th century.

Since CPC codes Y10S901 and Y10T436 belong to the ‘raccord’ class Y10 “Technical subjects covered by former USPC”, their in-hindsight re-assignment to older inventions before the CPC system was conceived is scattered and largely missing. For this reason, we discard these codes for the remainder of the analysis. The aforementioned clustering exercise is also supported by correlation analysis; a heatmap representation of the correlation matrix between the relevant series is reported in [Appendix B](#).

The evidence presented so far suggests, first, that given the emergence of strong heterogeneity among the underlying knowledge base encompassed by LS robotic artefacts, robots and intelligent machines are somewhat more than a simple GPT; second, that the time evolution of underlying CPC codes is rather distinctive; third, that clusters of innovations, comovements and anti-comovements appear, corresponding to successive industrial revolutions. This evidence is still however inconclusive about the emergence of recurrent long waves of innovations or of rather erratic technological clustering. The next [section](#) advances our understanding in this respect.

3 Regular waves or erratic constellations?

In this last step of the analysis, we investigate whether the assignment of our target CPC codes exhibit oscillations over time, intended as repeated cycles of upswings and downswings or rather less regular innovation waves determined by comovements of some bundle of innovations. Silverberg (2007) provides a synthetic reappraisal of the literature detecting long waves in innovation cycles. Two strands of literature still contend the interpretation of waves in innovation, a debate partly hosted by the *Cambridge Journal of Economics* in 1980s: on the one hand, the regular long waves perspective, following the Kondratiev (1935) fifty years cycle, sustained by Mensch (1979), who reports innovation clusters in recessionary phases, and only partly supported by Kleinknecht (1990); on the other hand, an anti-regular cycle perspective which, although denying the pure random walk hypothesis put forward by Solomou (1986), sustains the presence of long-cycles, although irregular, regulated by socio-economic technical conditions (Clark et al., 1981).

Criticisms on the robustness of the findings entail both the nature of the data used to measure innovative activities, wherein a lot of arbitrariness arises to account for basic innovations (see the criticism to Mensch, 1979), and the use of short-time periods actually impeding the genuine detection of long waves. In the following, relying on a long timespan dataset, collecting only potential *breakthrough* innovations, we attempt at overcoming these criticisms.

3.1 Wavelet analysis

In order to detect the presence of long waves, signal processing is an appropriate tool of investigation. Periodic oscillations are typically revealed by Fourier analysis: therein, the Fourier transform decomposes a signal into its constituent frequencies, thereby translating




Figure 5: Representation of a real-valued Morlet wavelet with $\omega_0 = 6$.

it from *time* domain to *frequency* domain. When applied to long time series, however, the Fourier approach only captures periodic behaviours that are detectable throughout the whole time frame with constant wavelength, while it fails to account for shorter periods of wavy dynamics and oscillations with time-varying wavelength. In simple terms, with the Fourier transform, the time information carried by the signal is lost. Moreover, the Fourier transform is known to be unsuitable for analysing non-stationary time series presenting irregular behaviours (Charpe et al., 2019; Gallegati, 2019).

In order to overcome this limitation, a 2-dimensional technique, which simultaneously accounts for the time *and* frequency domain, is needed. Wavelet analysis decomposes the original signal into a complete time-frequency representation, thereby retaining all the relevant information carried by the signal (for an introduction with applications in economics, see Aguiar-Conraria and Soares, 2014). Strictly speaking, the wavelet transform corresponds to a *convolution* of a time function (the time series under analysis) and a set of *wavelets*, namely wave-like functions satisfying certain properties and acting as the underlying *basis*. A more formal definition of relevant concepts is provided in Appendix C.

We deem this method superior to the short-time Fourier transform, another widely used 2-dimensional technique which splits a time signal into short segments and then computes the Fourier transform separately on each shorter segment, since the latter requires an additional tuning in the choice of the windowing function and related segment length, which by construction is fixed and therefore bounds the frequency resolution from below. Under the wavelet approach, instead, the time-frequency window is adaptive and local, and strikes a better balance between time domain and frequency domain. When it comes to the choice of the ‘mother’ wavelet function, we opt for the classical Morlet wavelet, depicted in Fig. 5 and formally defined in Appendix C. This is a common choice for generic data which is not ex-ante known to display a certain shape (such as electrocardiogram, to name one) and proves to have desirable properties, the explanation of which goes beyond the scope of this paper.⁷

Fig. 6 shows the time-frequency representation of the wavelet power spectrum (also known as the *scaleogram*) for each of our target CPC codes, obtained by applying the continuous wavelet transform to the original series (cf. Fig. 2). The vertical axis denotes the

⁷In particular, the Morlet wavelet has optimal joint time-frequency concentration and minimum possible uncertainty of the corresponding Heisenberg box, compared with other wavelet functions (Theolis, 1964).




Figure 6: Time-frequency representation of the wavelet power spectrum of target CPC series. Warmer (colder) coloured pixels represent higher (lower) values of the underlying coefficients. White contour lines denote significant areas at the 95% confidence level, against the null hypothesis of a white-noise model. Black points/lines represent local ridges of wavelet power.

oscillation period in years (defined as the reciprocal of frequency) in logarithmic scale, while the horizontal axis denotes the time location of the oscillation throughout our reference time frame. Warmer (respectively, colder) coloured pixels represent higher (lower) values of the underlying coefficients, meaning a more (less) pronounced periodic behaviour at the corresponding time-frequency coordinates. Shaded regions at the beginning and end of the period bound the so-called *cone of influence*, which excludes areas where edge effects might bias the analysis (see e.g. Aguiar-Conraria and Soares, 2014). White contour lines denote significant areas at the 95% confidence level, against the null hypothesis of a white-noise model. Black points/lines represent local ridges of wavelet power.

With respect to our analysis, the scaleograms synthetically provide three types of information. First, the very existence of periodic behaviour in a 3-digit CPC, a rather coarse level of aggregation, is informative of an underlying co-occurrence of LS technological artefacts, e.g. heterogeneous innovations within each technological class, exhibiting synchronised upsurges and subsequent declines (with colour temperature signalling the intensity of this coordinated movement). Second, the clustering length of innovations, or let us say, the innovation cycle, can be read in the vertical axis. The maximum oscillation period is set to 64 years, in accordance with the long wave theory.⁸ Third, the timing of occurrence of any potential wave is presented in the horizontal axis, in order to identify the underlying beginning and end dates. Therefore, predominant Kondratiev waves, if any, will be displayed as warm coloured spots located at the top of the panel. In general, the shorter the span of an innovation cycle, the lower a warm coloured spot will be located, while the less intense the coordinated upswing and downswing, the colder its colour. Note that the length of our dataset (184 years) is as such that we do not incur in false-positive waves for the oldest innovations, while this might be the case for newer innovations. We shall address this issue later on.

Three different patterns of periodic behaviour emerge. Starting with CPC code B01, the warmer areas are concentrated both to the left- and right-hand sides of the quadrant, divided by a colder central area. Therefore, two alternate phases of innovations cycles are detectable: an early phase with waves ranging from 16 up to 32 years ending around the 1900s and a second phase emerging around the 1920s, with longer periodic oscillations (warmer colours in the top-right corner). Additionally, shorter cycles, with lengths between 2-16 years (so-called Kitchin and Juglar cycles) are present in the yearly phase up to 1850 as well. This highlights the presence of innovative activities in “Physical or chemical processes or apparatus in general” (take the case, for example, of the du Pont family) which is known to mark the U.S. Industrial Revolution. The same pattern also characterises codes G01 and G05. Indeed, the latter CPC codes belong to the third innovation cluster, labelled as plateaued technologies, with functions like testing, control, measuring, and applications in physical or chemical processes. These technologies are indeed intimately related to both the first U.S. Industrial Revolution, characterised by textile and railways, and the third ICT phase.

⁸Also note that 64 years, which amounts to roughly one third of the overall length of our series, is a plausible cut-off for identifying meaningful waves using wavelet decomposition.

Long Kondratiev waves	Dominant CPCs	Coexisting CPCs
1780–1840: MECHANISATION AND TEXTILE	B23	A61, C12, G02, G06
1840–1890: STEAM POWER AND RAILWAYS	B01, G01, G05	B25, B62, B65
1890–1940: ELECTRICAL AND ENGINEERING	H01, H04	B25, B62, B65
1940–1990: MASS PRODUCTION AND AUTOMOTIVE	B25, B62, B65	G06, C12
1980– <i>ongoing</i> : ICT	C12, G02, G06	G01, G05, H01, H04

Table 2: Correspondence between target CPC codes and technological systems identified by Freeman and Louçã (2001).

The second type of behaviour is the one characterising codes B23, B25, B62, B65. These codes belong to the second innovation cluster, labelled hump-shaped, or equivalently, saturated technologies. Indeed the scaleogram shows a larger concentration of the signal both in the bottom-left quadrant and/or in the middle area. The cluster is indeed heterogeneous, with e.g. machine tools and metal working (B23) which show a statistically significant wave of approximately 16 years up to 1875, consistently with the mechanisation and textile phase. Innovations regarding the mass production phase, involving the assembly line and manual tools (conveying, packing, storing) in the automotive sector (B65) show short-term investment cycles between 1925 and 1975, but also long-lasting Kuznets cycles starting around 1850. Scaleograms are also useful for detecting transients in the underlying frequencies. This is the case of CPC code B25, comprising hand tools and manipulators, starting with a long period component of around 60 years (warmer area in the top-left corner) and then gaining frequency (and momentum) at the turn of the century, giving rise to significant 30–35 year long waves characterising the dynamics of the first three quarters of the century. This behaviour is also clearly visible to the naked eye in the corresponding panel of Fig. 2.

The third type of behaviour characterises codes A61, C12, G02, G06, H01, and H04, all showing warmer upper-right corners. Notably, they belong to our second cluster (ever increasing assignment) whose series exhibit a steadily growing trend (cf. Fig. 4(b)). While this cluster stands at the core of the ICT revolution, the wavelet transform is tricked into believing that the whole dynamics, and especially the steepest chunks therein, constitute an increasing portion of a very long period oscillation. It is not surprising, therefore, that the scaleograms all have their maximum coefficients concentrated in the upper-right corner. Indeed, the ever increasing trend starting in the 1960s obscures previous dynamics.

In order to obtain a more telling picture, we compute the wavelet power spectrum on a shorter segment of the original series, i.e. between 1836 (the beginning of our series) and 1950, corresponding to a period of more stationary dynamics, and consequently reducing the maximum oscillation period to 32 years. The new scaleograms, shown in Fig. 7, appear more insightful than their old counterpart, and highlight some early periodic behaviour that was previously invisible. Significant 32-year waves, ending around 1875, characterise codes A61, C12, G06, while lower cycles in the early phase also characterise G02.

Overall, our wavelet analysis is broadly in line with the long-wave perspective by Freeman and Louçã (2001). Each periodic episode identified therein maps to some domin-




Figure 7: Time-frequency representation of the wavelet power spectrum for series in the 2nd cluster. Warmer (colder) coloured pixels represent higher (lower) values of the underlying coefficients. White contour lines denote significant areas at the 95% confidence level, against the null hypothesis of a white-noise model. Black points/lines represent local ridges of wavelet power.

ant technological paradigms and key sectors of activity. As shown in Table 2, the target CPC codes retrieved from currently LS robotic patents allow to properly characterise successive waves of radical technological clustering. In this respect, our investigation offers new methodological underpinnings to study the long-term history of technological development, as such characterised by LS technologies, and related heuristics for its whole unfolding.

However, our results partly dissent from the periodicity and exact timing approach. Indeed, some dating discrepancies arise because of the covered period: since our dataset is based on U.S. inventions starting in the 1830s, the very first mechanisation wave is only partly captured.

Regarding remaining episodes, target CPC codes can be easily matched with dominant sectors/operations/functions identified by Freeman and Louçã (2001) within each wave. Nevertheless, first, we are not able to identify general Kondratiev waves lasting 50 or more years for all technologies; second, there is no upswing or downswing phase neatly distinguishing one wave from its predecessor or successor. If any, our results better fit with the ‘periodicisation’ of the three industrial revolutions, of which we are able to capture the declining phase of mechanisation, the upswing and downswing of steam power, and the upsurge of electricity, chemical and computing.

Provided that the waves we identified are *erratic, irregular* in length, and mostly *coexisting*, not only with other technologies within each cluster, but also between clusters, as shown by the presence of coexisting target CPC codes (cf. Table 2), we deem the more nuanced *development/system block* approach put forward by Nuvolari (2019) more appropriate to interpret our findings. The latter integrates the long wave perspective with the industrial revolutions dating, emphasising the autocatalytic processes behind these development blocks.

3.2 Comovements with GDP growth

Within the debate on the occurrence of long waves, a still unresolved issue pertains to the source of their emergence. The perspective of Mensch (1979) on radical innovations, and eventually resulting long waves, considers the latter to be more concentrated in periods of recessions, representing a managerial and organisational response to reduced market opportunities. If this is true for general innovations, the argument and related evidence should be even stronger with reference to LS ones, which are by their inner heuristics devoted to cost-cutting scopes and increased room for efficiency. Against the recession-driven perspective, Clark et al. (1981) consider radical innovations appearing in a rather heterogeneous unpredictable manner, both across sectors and time. Conversely, Keynesian driven approaches postulate innovative ideas clustering during economic expansions, because of markets opportunity and sales dynamics. A similar conclusion is reached in the demand-pull approach proposed by Schmookler (1962), who records inventive activities lagging behind sales dynamics.

In the following, we shall address the possible emergence of recession- or, alternatively, expansion-driven innovative propensity towards LS technologies. We start by investigating the relationship between our variable of interest, namely the family of target CPC codes signalling more automation/mechanisation oriented innovations, and GDP growth. Historical GDP data for the U.S. is retrieved from the Maddison Project Database, version 2018⁹ (Bolt et al., 2018). We distinguish LS innovative activity according to the three clusters identified in Fig. 4; each cluster corresponds to the sum of its underlying normalised series. Fig. 8 presents the relevant scatter plots, with GDP growth on the horizontal axis and CPC assignment on the vertical axis. In all three cases there is no meaningful correlation structure. However, it is possible to notice that highest levels of innovative activity are typically reached during periods of moderate growth, below 5% year on year. The absence of a clear correlation suggests that there is no such thing as an external push towards automation and mechanisation, but rather that the latter belong to a *natural trajectory* in the sense of Nelson and Winter (1982).

One may also argue that the absence of correlation derives from the inherent different dynamics of GDP growth, exhibiting volatility clustering, and innovation propensity in mechanisation/automation, none of whose clusters being stationary. Moreover, comovements may occur at lower frequencies. To this purpose, we further apply a Christiano-Fitzgerald (CF) band-pass filter (see Christiano and Fitzgerald, 2003) to all the involved time series, with the aim at separating the cycle component from the trend component. We parametrise the filter to retain oscillations with a minimum period of 30 years and a maximum period of 120 years, in order to keep very long-term information about the signal. Far from willing to accurately date the five waves in Table 2 (as in Gallegati, 2019)¹⁰, our aim is to spot the presence of any long-term comovement occurring between the two aggregates, when only the long-term signal is retained.

Fig. 9 presents the CF-filtered series superimposed to original unfiltered data regarding

⁹<https://www.rug.nl/ggdc/historicaldevelopment/maddison/releases/maddison-project-database-2018>

¹⁰Note that the number of waves is very sensible to the length of the cycle component.

innovation clusters. Comovements strongly differ from cluster to cluster. Overall, while the first and third clusters show a higher degree of lagged synchronisation, the second one has a complete autonomous pattern with respect to aggregate economic performance. Regarding the first cluster, the so-called hump-shaped technologies, panel (a) shows that long-term waves of economic growth generally precede mechanisation/automation waves. This is particularly the case for the initial boom around 1850, for the Great Depression, and the post-war economic boom. Nonetheless, the overall lag between peaks and troughs is rather long, at around ten years. A similar lagged comovement behaviour is also present in panel (e), regarding the third cluster, the so-called plateaued technologies. Again, whenever there are comovements, long lags emerge, and the GDP precedes the innovative wave. Additionally, peaks and troughs exhibit rather differing amplitudes. A clearly unrelated trend is the one in panel (c). The underlying cluster, characterised by an ever-increasing trend, presents a dynamics which is hardly reconcilable with economic activity in terms of amplitude, peaks and troughs. This cluster, which accounts for technologies linked to the electrification phase (first wave up to 1980s) and the following ICT revolution (second wave starting in the eighties), is the still dominant paradigm, relatively unaltered even after the Great Recession.

Given the lack of a unique comovement behaviour, we proceed by analysing (bivariate) wavelet cross power spectra (formally defined in [Appendix C](#)), in order to refine our visual inspection and detect otherwise under-looked patterns. The cross wavelet power spectrum allows to capture the underlying covariance structure between two time series in both time and frequency domains. With respect to the standard cross-correlation analysis, bivariate wavelet analysis permits the identification of not only leading and lagging indicators, but also of phase differences occurring at changing frequencies. Panels (b), (d), and (f) of Fig. 9 present the three power spectra computed on CF-filtered series. Arrows pointing to the right (respectively, left) indicate that GDP growth and LS innovation are in-phase (out of phase) at the respective time-frequency coordinates. Arrows pointing up-right and down-left signal that GDP locally leads innovation, while arrows pointing up-left and down-right suggest the opposite relation.

In all three cases, the cross wavelet power spectrum shows that the covariance is concentrated in periods between 32 and 64 years. Out-of-phase dynamics is present in the first cluster (panel (b)), whereby, contrary to previous visual inspection, the left-upward horizontal arrows concentrated around the 32-year period indicate innovative activities leading GDP. The leading structure persists for about a century, starting approximately in 1875 and ending in 1975. Considering that underlying CPC codes are primarily linked to mechanisation and automation, the leading structure suggests that LS activities negatively correlate with GDP growth in the long run (from thirty years onward). The leading thirty-year structure of innovation can be detected in panel (a), noting that the first and second blue peaks (whose horizontal distance equals the considered wavelength) precede the second and third orange peaks, given that the first one lies outside the cone of influence. The same lagging structure characterises the other two blue peaks (third and fourth) and the corresponding orange ones (fourth and fifth). With increasing wavelength

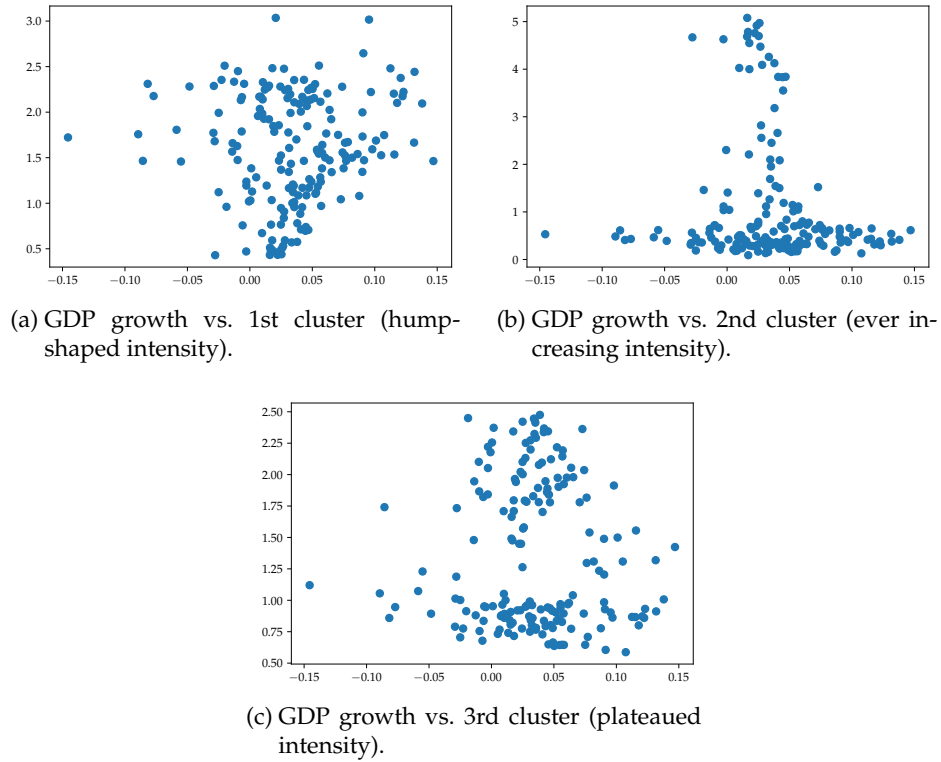


Figure 8: Scatter plots of GDP growth against LS patent intensity by innovation cluster.


(e.g. between the first and third peak), the phase structure disappears as shown by vertical arrows.

In-phase dynamics characterises both the second and third clusters with about 50-year wavelength. With respect to the second cluster (panel (d)), the in-phase dynamics is driven by the two subsequent troughs of GDP and innovative activities, while regarding the third cluster (panel (f)), it is driven by the second and third blue peaks, which are almost synchronised with the orange ones (second and third, and the fourth respectively). The 50-year in-phase dynamics however coexists with multiple areas of absence of any phase-structure, displayed along shorter and longer periods, as per the vertical arrows, in line with our visual inspection of comovements.


Overall, our analysis confirms the Clark et al. (1981) perspective of an autonomous and independent propensity to innovate in LS artefacts. In order to suitably detect synchronisation and phase structure one has to go beyond the simple time domain analysis and investigate simultaneously the time-frequency domain. In the latter case, covariances emerge but appear along different frequencies and corresponding wavelengths. Indeed, deterministic patterns, such as invariant leading-lagging structures, are hardly detectable.

4 Discussion and conclusions


This paper, relying on a long-term, still relatively unexplored, dataset on U.S. patenting activity, provides empirical evidence on the history of automation innovation, back since




(a) CF-filtered GDP growth vs. 1st cluster (hump-shaped intensity) CF-filtered and unfiltered series.




(b) Cross power spectrum of CF-filtered GDP growth and CF-filtered 1st cluster.



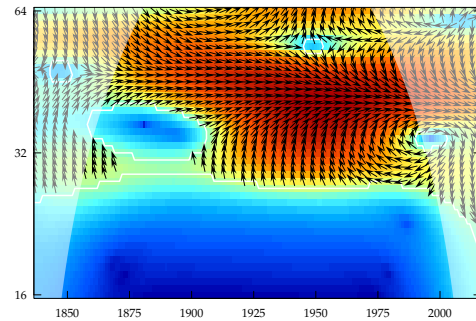
(c) CF-filtered GDP growth vs. 2nd cluster (ever increasing intensity) CF-filtered and unfiltered series.



(d) Cross power spectrum of CF-filtered GDP growth and CF-filtered 2nd cluster.



(e) CF-filtered GDP growth vs. 3rd cluster (plateaued intensity) CF-filtered and unfiltered series.



(f) Cross power spectrum of CF-filtered GDP growth and CF-filtered 3rd cluster.

Figure 9: Long-term comovements between GDP growth and LS patent intensity by innovation cluster. CF-filtered series of GDP growth (orange) and patent intensity (blue), and original patent intensity series (green) are pictured superimposed in panels (a), (c), and (e) for the 1st, 2nd, and 3rd clusters, respectively. Wavelet cross power spectra of CF-filtered GDP growth and patent intensity by cluster are pictured in panel (b), (d), and (f) for the 1st, 2nd, and 3rd clusters, respectively. Warmer (colder) coloured pixels represent higher (lower) values of the underlying coefficients. White contour lines denote significant areas at the 95% confidence level, against the null hypothesis of a white-noise model. Arrows denote phase difference and lead/lag relationships.

1830s. The labour-saving heuristics identified by Montobbio et al. (2020) via textual analysis on current robotic technologies allow to focus on a coherent set of technological CPC classes, the historical evolution of which is analysed in terms of timing, clustering, periodic behaviour, and comovements with GDP growth. The very fact that labour-saving CPCs differ widely in their assignment dynamics challenges the so-called GPT approach postulating a unique dominant technology, while it brings support to the idea of innovation waves seen as technological constellations.

Our findings are as follows. First, mechanisation and automation, or equivalently labour-saving heuristics, seem to constitute a “natural trajectory” (Nelson and Winter, 1982) in the evolution of the capitalist system, rather than a recurrent pattern. As opposed to socio-deterministic approaches linking the upsurge of automation to contingent phases, the hypothesis of natural trajectory implies that innovative efforts in labour-saving automation act as a background collective meta-heuristics, independent of local-focussing devices operating at the individual or firm level. Nonetheless, this notion does not exclude the formation of clustering patterns of innovation. Indeed, nothing pre-empts the coexistence of a Marxian interpretation of technical change, intended to mechanise and substitute labour to increase forms of control and appropriation over the production process by capitalists, by codifying into inanimate power previously non-codified knowledge, on the one hand, and the endogenous emergence of innovative efforts concentrated over a set of technological artefacts, in the Schumpeterian sense, on the other. There exist periods of more coordinated innovative effort resulting in upsurges and subsequent declines, highlighting some degree of technological clustering. In particular, we detect the presence of three technological clusters exhibiting distinct temporal patterns: from hump-shaped, to plateaued, to ever-increasing dynamics. Overall, the tension between invariant patterns, as expressed by socio-economic meta-routines granted by institutions and meta-institutions (Dosi et al., 2020), and unfolding heterogeneities characterising historical episodes and sectors of activity (Capone et al., 2019), remains largely unresolved in our interpretation of the socio-economic fabric.

Second, in detecting the presence of periodic behaviour via wavelet analysis, we do not confirm the presence of 50-year long Kondratiev waves. Indeed, we are not able to identify in mechanisation and automation subsequent regular periodic waves, leading to new technological systems. However, the dominant CPC codes characterising erratic technological constellations are in line with the technological system dating proposed by Freeman and Louçã (2001). The two seemingly contrasting results are instead coherent with the system block approach and the coexistence of both within-paradigm and between-paradigm trajectories (Dosi, 1982).

Third, with respect to the existence of a relationship between economic performance and innovative activity in mechanisation and automation, we hardly find any evidence thereof. Neither recession-driven nor euphoria-driven innovations are found, given the absence of correlation between GDP growth and patenting activity. When looking at the long-term cycle component, innovation and GDP growth present delinked patterns of waves, with heterogeneous troughs and peaks. Whenever comovements occur, waves

in GDP growth seem to precede, rather than follow, technological innovations, although the picture gets more nuanced when looking at both time and frequency domains together. In this regard, any purported saturation of the technological frontier or of innovative ideas are not detectable from the trends in innovation directed at the mechanisation and automation of tasks. Labour-saving efforts are there and involve a large set of technological artefacts, producers, and sectors of activity (Montobbio et al., 2020). This occurs rather independently of economic cycles at the macro-level. Our results however are not intended to deny the existence of local and discrete focussing devices or search heuristics which guide the innovative process. For instance, recent micro-evidence highlights the role of bottlenecks and opportunities shaping the innovation trajectory in Swedish manufacturing (Taalbi, 2017). Indeed, technological trajectories remain locally bounded by technological bottlenecks and market opportunities. Additionally, there might be other variables, such as the degree of union power, wage levels, and conflictual attitudes against mechanisation/automation moves, in line with the socio-technical approach (Noble, 1986), which might trigger the innovation dynamics. Conflictual claims about labour conditions might however result also in new technology meant to increase ergonomics and safety conditions in the workplace. The direction is therefore not univocal and the aforementioned relationships may represent avenues of future research.

The main limitation of our results comes from the level of aggregation: indeed, 3-digit CPC codes are rather heterogeneous and might also include labour-friendly innovations, even in their conception phase. Additionally, whether an innovation is labour-saving or labour-friendly is a question that pertains to the use of the artefact and its implementation in the production and organisational processes occurring at the firm and sectoral level. In the present paper, we look at the direction of innovative ideas, without reaching a conclusive appraisal on the employment effect from their use. Finally, given the widely heterogeneous and complex nature of technology, our investigation focusses on labour-saving innovations uniquely derived by current robotic artefacts, therefore potentially neglecting other labour-saving innovations sprung by different artefacts, not specifically linked to robotic automation. Future research would benefit from encompassing a wider investigation across the whole set of patents.

References

- Acemoglu, Daron and Pascual Restrepo (2018). “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment”. *American Economic Review* 108(6), pp. 1488–1542. DOI: [10.1257/aer.20160696](https://doi.org/10.1257/aer.20160696).
- Acemoglu, Daron and Pascual Restrepo (2019). “Automation and New Tasks: How Technology Displaces and Reinstates Labor”. *Journal of Economic Perspectives* 33(2), pp. 3–30. DOI: [10.1257/jep.33.2.3](https://doi.org/10.1257/jep.33.2.3).
- Acemoglu, Daron and Pascual Restrepo (2020). “Robots and Jobs: Evidence from US Labor Markets”. *Journal of Political Economy* 128(6), pp. 2188–2244. DOI: [10.1086/705716](https://doi.org/10.1086/705716).

- Aguiar-Conraria, Luís and Maria Joana Soares (2014). "The continuous wavelet transform: moving beyond uni- and bivariate analysis". *Journal of Economic Surveys* 28(2), pp. 344–375. DOI: [10.1111/joes.12012](https://doi.org/10.1111/joes.12012).
- Atack, Jeremy, Robert A. Margo and Paul Rhode (2020). "Mechanization Takes Command: Inanimate Power and Labor Productivity in Late Nineteenth Century American Manufacturing". NBER Working Paper No. 27436.
- Baines, Edward (1835). *History of the Cotton Manufacture in Great Britain*. H. Fisher, R. Fisher, and P. Jackson, London.
- Barbieri, Laura, Mariacristina Piva and Marco Vivarelli (2018). "R&D, embodied technological change, and employment: evidence from Italian microdata". *Industrial and Corporate Change* 28(1), pp. 203–218. DOI: [10.1093/icc/dty001](https://doi.org/10.1093/icc/dty001).
- Bolt, Jutta, Robert Inklaar, Herman DE Jong and Jan Luiten van Zanden (2018). "Rebasing 'Maddison': new income comparisons and the shape of long-run economic development". Maddison Project Working paper 10.
- Calvino, Flavio and M. Enrica Virgillito (2018). "The innovation employment nexus: A critical survey of theory and empirics". *Journal of Economic Surveys* 32(1), pp. 83–117. DOI: [10.1111/joes.12190](https://doi.org/10.1111/joes.12190).
- Capone, Gianluca, Franco Malerba, Richard R. Nelson, Luigi Orsenigo and Sidney G. Winter (2019). "History friendly models: retrospective and future perspectives". *Eurasian Business Review* 9(1), pp. 1–23. DOI: [10.1007/s40821-019-00121-0](https://doi.org/10.1007/s40821-019-00121-0).
- Castaldi, Carolina, Roberto Fontana and Alessandro Nuvolari (2009). "Chariots of fire: the evolution of tank technology, 1915–1945". *Journal of Evolutionary Economics* 19(4), pp. 545–566. DOI: [10.1007/s00191-009-0141-0](https://doi.org/10.1007/s00191-009-0141-0).
- Charpe, Matthieu, Slim Bridji and Peter Mcadam (2019). "Labor share and growth in the long run". *Macroeconomic Dynamics*, pp. 1–38. DOI: [10.1017/S1365100518001025](https://doi.org/10.1017/S1365100518001025).
- Christiano, Lawrence J. and Terry J. Fitzgerald (2003). "The band pass filter". *International economic review* 44(2), pp. 435–465. DOI: [10.1111/1468-2354.t01-1-00076](https://doi.org/10.1111/1468-2354.t01-1-00076).
- Cirillo, Valeria, Matteo Rinaldini, Jacopo Staccioli and Maria Enrica Virgillito (2020). "Technology vs. workers: the case of Italy's Industry 4.0 factories". *Structural Change and Economic Dynamics*. In press.
- Clark, John, Christopher Freeman and Luc Soete (1981). "Long waves and technological developments in the 20th century". In: *Konjunktur, Krise, Gesellschaft*. Ed. by D. Petzina and G. VAN Roon. Klett-Cotta, Stuttgart.
- Cohen, Michael D., Roger Burkhart, Giovanni Dosi, Massimo Egidi, Luigi Marengo, Massimo Warglien and Sidney Winter (1996). "Routines and other recurring action patterns of organizations: contemporary research issues". *Industrial and Corporate Change* 5(3), pp. 653–698. DOI: [10.1093/icc/5.3.653](https://doi.org/10.1093/icc/5.3.653).

- Dahmén, Erik (1988). “‘Development blocks’ in industrial economics”. *Scandinavian Economic History Review* 36(1), pp. 3–14. DOI: [10.1080/03585522.1988.10408102](https://doi.org/10.1080/03585522.1988.10408102).
- Dosi, Giovanni (1982). “Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change”. *Research Policy* 11(3), pp. 147–162. DOI: [10.1016/0048-7333\(82\)90016-6](https://doi.org/10.1016/0048-7333(82)90016-6).
- Dosi, Giovanni (1988). “Sources, procedures, and microeconomic effects of innovation”. *Journal of Economic Literature* 26(3), pp. 1120–1171.
- Dosi, Giovanni, Luigi Marengo and Alessandro Nuvolari (2020). “Institutions and economic change: some notes on self-organization, power and learning in human organizations”. *Eurasian Business Review* 10, pp. 1–22. DOI: [10.1007/s40821-019-00132-x](https://doi.org/10.1007/s40821-019-00132-x).
- Dosi, Giovanni, Mariacristina Piva, M. Enrica Virgillito and Marco Vivarelli (2019). “Embodied and Disembodied Technological Change: The Sectoral Patterns of Job-Creation and Job-Destruction”. IZA Discussion Paper No. 12408.
- Freeman, Chris (2019). “History, co-evolution and economic growth”. *Industrial and Corporate Change* 28(1), pp. 1–44. DOI: [10.1093/icc/dty075](https://doi.org/10.1093/icc/dty075).
- Freeman, Christopher and Francisco Louçã (2001). *As time goes by: from the industrial revolutions to the information revolution*. Oxford University Press.
- Freeman, Christopher and Luc Soete (1987). *Technical change and full employment*. Wiley-Blackwell.
- Gallegati, Marco (2019). “A system for dating long wave phases in economic development”. *Journal of Evolutionary Economics* 29(3), pp. 803–822. DOI: [10.1007/s00191-019-00622-1](https://doi.org/10.1007/s00191-019-00622-1).
- Graetz, Georg and Guy Michaels (2018). “Robots at Work”. *Review of Economics and Statistics* 100(5), pp. 753–768. DOI: [10.1162/rest_a_00754](https://doi.org/10.1162/rest_a_00754).
- Harrison, Rupert, Jordi Jaumandreu, Jacques Mairessee and Bettina Peters (2014). “Does innovation stimulate employment? A firm-level analysis using comparable micro-data from four European countries”. *International Journal of Industrial Organization* 35, pp. 29–43. DOI: [10.1016/j.ijindorg.2014.06.001](https://doi.org/10.1016/j.ijindorg.2014.06.001).
- Kleinknecht, Alfred (1990). “Are there Schumpeterian waves of innovations?” *Cambridge Journal of Economics* 14(1), pp. 81–92. DOI: [10.1093/oxfordjournals.cje.a035120](https://doi.org/10.1093/oxfordjournals.cje.a035120).
- Kondratiev, Nikolai D. (1935). “The long waves in economic life”. *Review of Economics and Statistics* 17(6), pp. 105–115. DOI: [10.2307/1928486](https://doi.org/10.2307/1928486).
- Lachenmaier, Stefan and Horst Rottmann (2011). “Effects of innovation on employment: A dynamic panel analysis”. *International Journal of Industrial Organization* 29(2), pp. 210–220. DOI: [10.1016/j.ijindorg.2010.05.004](https://doi.org/10.1016/j.ijindorg.2010.05.004).

- Martinelli, Arianna (2012). "An emerging paradigm or just another trajectory? Understanding the nature of technological changes using engineering heuristics in the telecommunications switching industry". *Research Policy* 41(2), pp. 414–429. DOI: [10.1016/j.respol.2011.10.012](https://doi.org/10.1016/j.respol.2011.10.012).
- Marx, Karl (1955). *The Poverty of Philosophy*. Progress Publishers, Moscow.
- Mensch, Gerhard (1979). *Stalemate in technology: innovations overcome the depression*. Ballinger Pub Co.
- Montobbio, Fabio, Jacopo Staccioli, Maria Enrica Virgillito and Marco Vivarelli (2020). "Robots and the origin of their labour-saving impact". LEM Working Paper Series 2020/03.
- Moro, Angelo, Matteo Rinaldini, Jacopo Staccioli and M. Enrica Virgillito (2019). "Control in the era of surveillance capitalism: an empirical investigation of Italian Industry 4.0 factories". *Journal of Industrial and Business Economics* 46(3), pp. 347–360. DOI: [10.1007/s40812-019-00120-2](https://doi.org/10.1007/s40812-019-00120-2).
- Nelson, Richard R. and Sidney Winter (1982). *An evolutionary theory of economic change*. Harvard University Press.
- Noble, David (1986). *Forces of production: A social history of industrial automation*. Oxford University Press.
- Nuvolari, Alessandro (2019). "Understanding successive industrial revolutions: A 'development block' approach". *Environmental Innovation and Societal Transitions* 32, pp. 33–44. DOI: [10.1016/j.eist.2018.11.002](https://doi.org/10.1016/j.eist.2018.11.002).
- Pellegrino, Gabriele, Mariacristina Piva and Marco Vivarelli (2019). "Beyond R&D: the role of embodied technological change in affecting employment". *Journal of Evolutionary Economics* 29(4), pp. 1151–1171. DOI: [10.1007/s00191-019-00635-w](https://doi.org/10.1007/s00191-019-00635-w).
- Perez, Carlota (1983). "Structural change and assimilation of new technologies in the economic and social systems". *Futures* 15(5), pp. 357–375. DOI: [10.1016/0016-3287\(83\)90050-2](https://doi.org/10.1016/0016-3287(83)90050-2).
- Piva, Mariacristina and Marco Vivarelli (2018). "Technological change and employment: is Europe ready for the challenge?" *Eurasian Business Review* 8(1), pp. 13–32. DOI: [10.1007/s40821-017-0100-x](https://doi.org/10.1007/s40821-017-0100-x).
- Rosenberg, Nathan (1976). *Perspectives on Technology*. Cambridge University Press. DOI: [10.1017/cbo9780511561313](https://doi.org/10.1017/cbo9780511561313).
- Schmookler, Jacob (1962). "Economic sources of inventive activity". *Journal of Economic History* 22(1), pp. 1–20. DOI: [10.1017/s0022050700102311](https://doi.org/10.1017/s0022050700102311).
- Schumpeter, Joseph Alois (1939). *Business cycles*. Vol. 1. McGraw-Hill, New York.

- Silverberg, Gerald (2007). “Long Waves: Conceptual, Empirical and Modelling Issues”. In: *Elgar companion to neo-Schumpeterian economics*. Ed. by Horst Hanusch and Andreas Pyka. Edward Elgar Publishing. DOI: [10.4337/9781847207012.00060](https://doi.org/10.4337/9781847207012.00060).
- Silverberg, Gerald and Bart Verspagen (2003). “Breaking the waves: a Poisson regression approach to Schumpeterian clustering of basic innovations”. *Cambridge Journal of Economics* 27(5), pp. 671–693. DOI: [10.1093/cje/27.5.671](https://doi.org/10.1093/cje/27.5.671).
- Simonetti, Roberto, Karl Taylor and Marco Vivarelli (2000). “Modelling the employment impact of innovation”. In: *The Employment Impact of Innovation*. Ed. by Mario Pianta and Marco Vivarelli. Routledge. Chap. 3, pp. 26–43.
- Solomou, Solomos (1986). “Innovation clusters and Kondratieff long waves in economic growth”. *Cambridge Journal of Economics* 10(2), pp. 101–112. DOI: [10.1093/oxfordjournals.cje.a034987](https://doi.org/10.1093/oxfordjournals.cje.a034987).
- Taalbi, Josef (2017). “What drives innovation? Evidence from economic history”. *Research Policy* 46(8), pp. 1437–1453. DOI: [10.1016/j.respol.2017.06.007](https://doi.org/10.1016/j.respol.2017.06.007).
- Theolis, Anthony (1964). *Computational signal processing with wavelets*. Birkhäuser, Boston.
- Thompson, Edward P. (1963). *The Making of the English Working Class*. Victor Gollancz, London.
- Trajtenberg, Manuel (2019). “Artificial Intelligence as the Next GPT: A Political-Economy Perspective”. In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.
- Van Reenen, John (1997). “Employment and Technological Innovation: Evidence from U.K. Manufacturing Firms”. *Journal of Labor Economics* 15(2), pp. 255–284. DOI: [10.1086/209833](https://doi.org/10.1086/209833).
- Van Roy, Vincent, Dániel Vértesy and Marco Vivarelli (2018). “Technology and employment: Mass unemployment or job creation? Empirical evidence from European patenting firms”. *Research Policy* 47(9), pp. 1762–1776. DOI: [10.1016/j.respol.2018.06.008](https://doi.org/10.1016/j.respol.2018.06.008).
- Vivarelli, Marco (1995). *The Economics of Technology and Employment: Theory and Empirical Evidence*. Edward Elgar Pub.
- von Tunzelmann, G. Nick (1995). “Time-Saving Technical Change: The Cotton Industry in the English Industrial Revolution”. *Explorations in Economic History* 32(1), pp. 1–27. DOI: [10.1006/exeh.1995.1001](https://doi.org/10.1006/exeh.1995.1001).

Appendix A


Technical summary of Montobbio et al. (2020)

The contribution consists of three methodological steps. First, patents which either directly or indirectly relate to robotics technology are singled out. Second, a procedure is

implemented in order to detect the underlying LS heuristics and pinpoint the set of explicitly LS patents. Finally, a probabilistic topic model is estimated in order to devise a human-machine taxonomy.

Step 1 – Identification of robotic patents The analysis starts with the entire set of 3,557,435 patent applications published by the USPTO between 1st January 2009 and 31st December 2018. Robotic patents are pinpointed therein according to two distinct criteria, one based on the patent classification codes specified within applications, the other based on textual keyword search. A patent is deemed ‘robotic’ if it obeys at least one of the criteria. In particular, a robotic patent according to the first criterion (dubbed ‘CPC’) must be assigned by patent examiners at least one of a set of 174 full-digit CPC codes which reflect former U.S. Patent Classification (USPC) class 901 (“Robots”). Likewise, a robotic patent according to the second criterion (dubbed ‘K10’) must contain the word ‘robot’ in its full-text at least 10 times, including derivational and inflectional affixes. The first criterion identifies 10,929 robotic patents, while the second criterion identifies another 18,860 (after discarding robotic patents according to the first criterion). The two criteria single out a total of 29,789 robotic patents, i.e. approximately 0.84% of the original (universe) population.

Step 2 – Identification of labour-saving patents Labour-saving patents constitute a subset of robotic patents, identified by a multiple *word* co-occurrence query at the *sentence* level. In particular, a patent is deemed labour-saving (after an additional manual validation step) if its full-text contains at least one sentence in which the verbal predicate, direct object, and object attribute belong to the following lists:



In total, 1,276 labour-saving patents are found (approximately 4.3% of all robotic patents), of which 461 ($\approx 36.1\%$) belong to the CPC group and 815 ($\approx 63.9\%$) belong to the K10 group.

Step 3 – Probabilistic topic model and human-machine taxonomy The set of labour-saving patents is technologically characterised vis-à-vis the superset of robotic patents by leveraging the latent semantic structure of the whole collection of patents’ full-texts. The analysis proceeds along the following methodological workflow. First, a probabilistic topic model is estimated on the whole population of robotic patents, which associates a

distribution θ_d of membership over the K -dimensional set β of topics to each patent d . Second, a distribution of CPC codes, according to the original attribution of codes to each patent by the topic proportions θ_d found in the previous step, is associated to each topic β_k . Finally, the relevance of each topic to the whole population of robotic patents is compared to the same relevance to the subset of labour-saving patents, in order to draw quantitative conclusions on which technologies are relatively more and less relevant in characterising the two sets of patents. The probabilistic topic model, asked to identify $K = 20$ topics, returns each topic β_k as a list of relevant keywords and a membership value $\theta_{d,k}$ of each patent d to topic k . An aggregate measure of relevance of each topic k to an arbitrary collection of patents D (e.g. the set of robotic patents or labour-saving patents) can be defined as the simple average membership of all patents in the collection to topic k , as follows:

$$\Theta_k^D := \frac{\sum_{d \in D} \theta_{d,k}}{|D|} \quad \forall k = 1, \dots, K$$


When the underlying collection of patents D is the whole set of robotic patents, Θ_k^{rob} measures the relevance of each topic to robotic patents; analogously, when the underlying collection of patents D is the subset of labour-saving patents, Θ_k^{LS} measures the relevance of each topic to labour-saving patents. Distributions Θ_k^{rob} and Θ_k^{LS} are pictured in Fig. 1, where topics are sorted by decreasing relevance to the robotic patents collection. Finally, the relative synthetic measure $\tilde{\Theta}_k^{\text{LS}}$ mentioned in section 2 is defined as

$$\tilde{\Theta}_k^{\text{LS}} := \frac{\Theta_k^{\text{LS}}}{\Theta_k^{\text{rob}}} \quad \forall k = 1, \dots, K.$$

Appendix B

Correlation analysis

The following picture provides a heatmap representation of the correlation matrix of target CPC intensity series (CPC codes Y10S901 and Y10T436 are discarded).



At a first glance, it is immediate to spot a first cluster of codes, B23, B25, B62, and B65, whose assignment dynamics negatively correlates (lighter pixels) with all the other codes. At a second glance, a second cluster is present, with codes A61, C12, G02, G06, H01, and H04 exhibiting a highly positive correlation (darker pixels) with one another. The remaining codes, B01, G01, and G05, form a third cluster, in that they display milder, near-zero correlation coefficients with series of the second cluster, and negative correlation coefficients with series of the first.

Appendix C

Wavelet analysis definitions

Consider the Hilbert space $L^2(\mathbb{R})$ of square-integrable functions. A function $\psi(t) \in L^2(\mathbb{R})$ is called a *mother* wavelet if it satisfies the *admissibility* condition

$$\int_{-\infty}^{+\infty} \frac{|\Psi(\omega)|}{|\omega|} d\omega < +\infty, \quad (1)$$

where $\Psi(\omega)$ stands for the Fourier transform of $\psi(t)$.¹¹ Condition (1) implies that $\Psi(\omega)$ vanishes when frequency ω equals zero:

$$|\Psi(\omega)| \Big|_{\omega=0} = 0. \quad (2)$$

¹¹In practice, depending on the specific functional form of $\psi(t)$, additional *regularity* conditions may be required to ensure sufficient decay, beyond square integrability.

In other words, the wavelet must display a band-pass like spectrum. Moreover, condition (1) also requires that the wavelet's average value over time is zero:

$$\int_{-\infty}^{+\infty} \psi(t) dt = 0. \quad (3)$$

Loosely speaking, eqs. (2) and (3) together establish that $\psi(t)$ must wiggle up and down over time, and therefore resembles a wave-like function.

The Morlet wavelet, used in the analysis of this paper and pictured in Fig. 5, is composed of a complex exponential (carrier) multiplied by a Gaussian window (envelope):

$$\psi_{\omega_0}(t) = \pi^{1/4} \left(e^{i\omega_0 t} - e^{-\omega_0^2/2} \right) e^{-t^2/2}.$$

It is possible to show that if $\omega_0 \geq 5$, then $\psi_{\omega_0}(t)$ satisfies the admissibility condition (1). In the present analysis, we set $\omega_0 = 6$.

A (generic) mother wavelet $\psi(t)$ 'gives birth' to a family $\psi_{s,\tau}(t)$ of so-called *child* wavelets or wavelet *daughters* by means of scaling and translation operations

$$\psi_{s,\tau}(t) = \frac{1}{\sqrt{s}} \psi \left(\frac{t - \tau}{s} \right), \quad (4)$$

where $s \in \mathbb{R}^+$ denotes the scaling factor, which stretches/shrinks the mother wavelet, and $\tau \in \mathbb{R}$ denotes the translation parameter, which shifts the mother wavelet across time.

The continuous wavelet transform (CWT) of a function of time (or time series) $f(t)$, with respect to a mother wavelet $\psi(t)$, is defined as

$$W_{f,\psi}(s, \tau) = \frac{1}{\sqrt{|s|}} \int_{-\infty}^{+\infty} f(t) \psi^* \left(\frac{t - \tau}{s} \right) dt, \quad (5)$$

where the * superscript denotes complex conjugation. The CWT provides a representation of $f(t)$ in terms of wavelet basis functions $\psi_{s,\tau}(t)$, by letting the scale and translation parameters vary continuously. In other words, the CWT is a convolution of the signal $f(t)$ with the family of stretched and translated child wavelets defined in eq. (4).

Given eq. (5), the wavelet power spectrum, represented by a heatmap (scaleogram) in Figs. 6 and 7, is computed as

$$WPS_{f,\psi}(s, \tau) = |W_{f,\psi}(s, \tau)|^2.$$

In the bivariate case, the *cross* wavelet transform of two signals $f(t)$ and $g(t)$ with respect to a mother wavelet $\psi(t)$ is defined as

$$W_{f,g,\psi}(s, \tau) = W_{f,\psi}(s, \tau) \cdot W_{g,\psi}(s, \tau)^*,$$

while the *cross* power spectrum, pictured in panels (b), (d) and (f) of Fig. 9, is given by

$$XPS_{f,g,\psi}(s, \tau) = |W_{f,g,\psi}(s, \tau)|.$$