Humans, Robots and Values

Karen Renaud, Paul Cockshott[1,]

School of Computing Science University of Glasgow
[1,]

Abstract

The issue of machines replacing humans dates back to the dawn of industrialisation. In this paper we examine what is fundamental in the distinction between human and robotic work by reflecting on the work of the classical political economists and engineers. We examine the relationship between the ideas of machine work and human work on the part of Marx and Watt as well as their role in the creation of economic value. We examine the extent to which artificial power sources could feasibly substitute for human effort in their arguments. We go on to examine the differing views of Smith and Marx with respect to the economic effort contributed by animals and consider whether the philosophical distinction made between human and non-human work can be sustained in the light of modern biological research. We emphasise the non-universal character of animal work before going on to discuss the ideas of universal machines in Capek and Turing giving as a counter example a cloth-folding robot being developed in our School. We then return to Watt and discuss the development of thermodynamics and information theory. We show how recent research has led to a unification not only of these fields but also a unitary understanding of the labour process and the value-creation process.

We look at the implications of general robotisation for profitability and the future of capitalism. For this we draw on the work of von Neumann not only on computers but also in economics to point to the *real* threat posed by robots.

Email address: `william.cockshott@glasgow.ac.uk` (Karen Renaud, Paul Cockshott)

URL: `http://www.dcs.gla.ac.uk/~karen/` (Karen Renaud, Paul Cockshott)

1. Introduction

From the 1950s onwards the threat that automation posed to human labour became a persistent theme in popular science fiction [1, 2]. Authors explored what it meant to be human, by contrasting us with hypothetical robots. Such robots were generally seen as coming into existence centuries into the future. In the last decade the rate of progress in robotics has accelerated way beyond popular expectation. The timescales of Asimov and Dick look generous, whereas the dystopian near future of 'Player Piano' [3] seems grimly real. This anxiety is not limited to novelists. Even Stephen Hawking told the BBC:

> *"The development of full artificial intelligence could spell the end of the human race."* [4]

Robotics is made possible by advances in mechanical engineering but, above all, by informatics. In this essay we look at how ideas derived from informatics allow us a more precise view of what differentiates us from robots and, on the other hand, how information science can give us a deeper insight into the nature of human labour. Having gained this understanding, we can go on to examine what sort of threat robots really pose to us, as humans.

We commence by exploring the concept of labour and the differences between being being paid for labour, or the ability to labour, in Section 2. In Section 3 we consider alternative sources of value than labour, and conclude that, of all commodities, labour is the only one that matches best to prices. In Section 4 we explain the important distinctions between horses and humans, and explain why machines were indeed able to replace the former. Then in Section 5 we look at Turing's universal machine and the latest research into robot capabilities. In Section 6 we explore the links between the laws of thermodynamics and the concept of value and conclude, in Section 7, that human labour is still important. Section 8 revisits the seminal work of Von Neumann to consider, once again, the feasibility of a robot-dominated world. Section 9 reflects on the arguments made and Section 10 concludes.

2. Ideas of Work and Power

Marx famously made a distinction between labour and labour-power [5]. We will explain what he meant by this distinction.

Marx had a problem explaining the apparent dichotomy between these terms. On the one hand all market sales can be seen as fair and equal exchanges. On the other, the end result of these fair and equal exchanges was the production of something whereby one group of people became immensely wealthy at the expense of another group. How did this state of affairs result from a fair and equal exchange? One explanation could be that workers are cheated of the value of the labour: they are only paid part of the value of their labour because the market is rigged in such a way that they can never sell it for the full value.

Marx pondered how you could have a situation where it appears that the labourer is paid a fair price for his labour, which is the price, according to Ricardo [6], that is necessary to maintain and reproduce the labouring class, and at the same time there is profit and exploitation. And he, in effect, concludes: "Well, what is actually happening is that people are not being paid for their work, they are actually being paid for their ability to work."

A self-employed craftsman who makes something and sells it on the market, sells the product of his labour directly. Similarly with a roofer who comes and repairs your roof. They are paid directly for their actual work. If somebody is employed in a cotton mill to spin cotton, they do not sell the product of their labour; what they are selling is their *ability to labour*. The amount of labour that the employer can get out of a worker per day is a variable quantity. Its duration and intensity are variable. Characteristically of the time Marx was writing, working hours were extremely long, and had been progressively extended by the factory system. The intensity was, with mechanisation, tied to the speed at which the machinery in the mill operated. When the water was high in the river, the work was more intense:

> *O, dear me, the mill is running fast*
> *And we poor shifters canna get nae rest*
> *Shifting bobbins coarse and fine*
> *They fairly make you work for your ten and nine*
>
> *O, dear me, I wish this day were done*
> *Running up and doon the Pass is nae fun*
> *Shiftin', piecin', spinning warp, weft and twine*
> *To feed and clothe ma bairnie offa ten and nine*

The Jute Mill Song of the Dundee weavers, first recorded by Ewan MacColl, published by Mary Brooksbank [7]

What was being sold was the *ability to labour*. The amount of labour that the employer got out of that could be quite a lot more, according to the conditions of labour. This distinction between *labour* itself and *the ability to labour* must have some origin. Society must have prepared us for this distinction. So, where did these concepts actually come from?

There is reason to believe that this distinction originated at the start of the industrial revolution when Watt was producing steam engines. Watt didn't actually invent the steam engine. He was set the task of repairing a model Newcomb Engine, when he was a technician at Glasgow University. These engines were used for pumping: they had no rotary motion, because they were purely pumping engines. Watt looked at this, and because he had been working along with Black on the nature of heat [8], he realised that, in fact, these Newcomb engines were very inefficient because they threw away heat. They repeatedly cooled the piston down by condensing the steam in the piston by spraying water into it, and therefore a lot of the heat was wasted. What Watt actually did was to invent the separate condenser, whereby the heat was removed from the steam in a separate vessel. He also invented a series of automatic valves, which let the steam through from the piston into the condenser, or let steam into the other side of the piston, and this led to a great improvement in the efficiency of steam engines.

These steam engines were hired out by Watts' company, and he promised to hire them out for less than the saving in coal people would have made if they used a Newcomb Engine. In order to do that, he needed to have some way of measuring how much work these engines were doing and rating their power. Since the alternative to using an engine was to use horses, he became the first person systematically to study the amount of work a horse could do and thus introduce the concept of horsepower. In doing this he was the first person to make the distinction between the *ability to do work*, which was the power of the horse, and the *actual work*. For Watt, work was purely labour, pure effort: the effort of a horse or the effort of a person hauling up weights, physical exertion of effort. When you think of society in the late 18th century that is a very reasonable assessment of what work was, because most work entailed physical exertion of human muscle. That was primarily what people were being paid to do. Most of it was heavy, physical work.

In Adam Smith's writings this marrying of physical effort with labour

is such that Smith can talk with ease about a farmer having his labouring servants and his labouring cattle, because they are both seen as doing the same thing[2]. What Watt promises people is power, the ability to do work with his machines. Matthew Boulton, his partner, proudly announced to George II: *"Your Majesty, I have at my disposal what the whole world demands; something which will uplift civilization more than ever by relieving man of undignified drudgery. I have steam power"*. By this means he is going to transform the wealth of society. In a real sense he does this, because the power of his machines, within a few decades, are turning out more effort than all the muscle power of the human beings and horses in the kingdom. From that perspective he seems to have caught a key aspect of labour and of power, and that conceptualisation is still very much present in the classical political economists.

Smith says that labour is the original currency by which we win things from nature[3] and he also talks of labour as something that both humans and animals do. Both are seen, in that society, as labour. He knows that animals aren't skilled and that there is a limit to the labour they can do: they can not participate in the division of labour, for example. Smith is also interested in why things are valuable, and he quickly disposes of the idea that it is because they are useful. He points out that there are lots of valuable things that aren't particularly useful. The only constant is that things which *are* valuable require a lot of work to produce. At that point in time, at the dawn of the industrial age, when most labour was physical, the distinction between the kinds of labour that Watt and Smith were researching at Glasgow University, was not clear. They were both dealing with work: Smith was dealing with how work could be made more efficient by sub-dividing and specialising labour, and Watt was looking at how work could be replaced by

---

[2] "That part of the capital of the farmer which is employed in the instruments of agriculture is a fixed, that which is employed in the wages and maintenance of his labouring servants, is a circulating capital. He makes a profit of the one by keeping it in his own possession, and of the other by parting with it. The price or value of his labouring cattle is a fixed capital in the same manner as that of the instruments of husbandry. Their maintenance is a circulating capital in the same manner as that of the labouring servants" ([9] II.1.10)

[3] The annual labour of every nation is the fund which originally supplies it with all the necessaries and conveniences of life which it annually consumes, and which consist always either in the immediate produce of that labour, or in what is purchased with that produce from other nations. ([9]I.I.1)

artificial sources of power.

What is striking in Smith's 'Wealth of Nations' is that he does not discuss the use of powered machinery. All his economic improvements come from the sub-division of labour so that people can complete their tasks more quickly. By doing the same task again and again, they become more skilled at it, their movements become more automatic, and they do not lose time switching between tasks, and thereby more is produced. This vision of production is still pre-industrial, because powered industry didn't exist at that point. One stand-out exception was the water mills, but, apart from the production of flour, mass production wasn't generally based on powered machinery.

Marx takes the labour theory of value over from Smith and he makes it more precise in some ways because Smith confuses how much labour you can purchase with how much labour is required to make something, and he treats these as much the same thing. In a pre-industrial society of handicrafts and farmers they *are* essentially the same thing. When a Scottish farmer puts his grain on sale and buys, in return, some produce from the blacksmith, the value of his corn can be expressed in terms of how much of other people's labour he could command with it. He was indirectly commanding the labour of other tradesmen.

So, the the idea that value equals the amount of labour you can command has an intuitive appeal in a pre-capitalist, or only partially capitalist, society. Once capitalism becomes widespread it is not the same thing at all because wages only make up a part of the value of what is sold. Although something may require a certain amount of labour, the employer hasn't had to pay that much to his workers, so the item actually commands more labour than it requires to produce.

This distinction between labour *commanded* and labour *embodied* was pointed out by Ricardo, writing after the introduction of capitalist machinery in the late 19th Century [6]. These distinctions, which were not apparent in the mid-18th century, started becoming apparent in the early 19th Century, and Marx bases his distinction on the one Ricardo makes. He tries to explain how it is still possible that everything sells for its value and yet exploitation still results. His explanation is based on the distinction between power and work done, which we are arguing that he essentially derives from Watt.

3. Other Possible Sources of Value

Is this distinction between value and the power to create it specific to human labour? It is probably a property of whatever you take to be the value substance. Suppose you take the British economy, and instead of trying to calculate the labour value of all the products of the main industries, you choose to say: "we will treat oil as the substance of value", and "The value of each product is the number of barrels of oil that has to be directly or indirectly used to make that product." Now, for gasoline that's relatively simple. Plastics are also going to involve oil going into their production. But you then consider something like a TV set, and a TV set is going to require oil to have been burnt to provide the energy to transport it, oil being burnt to provide the energy to manufacture the parts, and then in addition it will have oil going into the plastics which make up the frame of it. You can systematically do this if you use the input/output tables: you can work out how much oil goes into everything. The obvious next logical question is: "What is the value of oil itself?" Well, if you are going to be consistent, you would say: "It is the amount of oil required to produce it." What do you discover?

You discover that it takes less than one barrel of oil, directly or indirectly, to produce a barrel of oil. If you took that definition, you would find that after one attempt at your calculations the value of a barrel oil is not one full barrel, but maybe a third of a barrel of oil has to be used to make a barrel of oil, so that the value of everything would have suddenly dropped by a third. The next time you attempt to calculate it, again it drops further. The point is that whatever commodity you take to be your value substance, in a functioning economy, it always requires less than one unit of the substance to make one unit of it. It is a property of saying something is the value substance, that you have to make that distinction.

You can do the same for steel, electricity, etc [10]. It is obvious that you have to take basic commodities, ones used in the production of everything else, directly or indirectly, to make the calculation. There is a limited number of these basic commodities: not all commodities are basic in that sense [11]. But, in what Marx called 'Sector 1' of the economy, these products are used directly or indirectly to make everything else. You could, in principle, choose any one of those and say: "this is the substance of value". You could get an internally consistent theory of value from that. The question is: would it be empirically accurate in predicting prices?

When you try investigating, you find that energy theories, oil theories or steel theories of value give a much poorer prediction of actual prices than the labour theory of value. Although, in principle, you could make that distinction for any commodity, it actually turns out that the only one that empirically makes sense is *labour*.

4. Humans and Horses

Since Watt's engines replaced horses as sources of power, why have other machines not replaced people as well? It comes down to the fact that there is more to work than just muscular energy. Where it is just muscular energy, it is not too difficult to replace with machinery. Thus horses got replaced very quickly, and somewhat later, the heavy labour of the navvies who dug the railways was replaced by the work of earth-moving machines and diggers of one sort or another. That only occurs when the wages rise enough for it to be worth using the machines, but once it is worth engaging machines, the people can indeed be replaced. The difference between a human being and a horse is that a horse can only pull: you can not set a horse to spin or tend the sails of a ship. Human beings, on the other hand, can be set to almost any labouring task.

They may not be terribly good at it, but they are adaptable, and it is this adaptability which must be a fundamental reason why we are the dominant species. We are, in effect, universal robots. And that, of course, is a bit of pun because the term robot in English comes from Capek's [12] play, which was written in the twenties and involves an English inventor Rossum, who comes to Prague and sets up a factory making mechanical men who are universal robots. And, of course, 'robot' in the Slavic language just means 'worker'. These fictional universal robots really did represent a threat to humans, because they could do any job that a human could do.

Such a universal robot would be a real existential threat to human beings because at that point the wealthy could see no point in maintaining a labouring population. In countries where they remained in control one hates to think what would happen. Bear in mind, and this is something that novelists have not been slow to point out, that if robots really were capable of replacing human beings in all these tasks, they would have to have as complex an internal life as we do. In Philip K. Dick's novel, "*Do Androids Dream of Electric Sheep*" [1], which was turned into the movie "*Blade Runner*", he

8

is essentially making the point that the Androids would rebel. The same capabilities would come with the same motivations.

The ideas that people form are obviously shaped by the technologies they grow up with [13]. The idea of a robot wasn't there in Marx's day. To our minds, the distinctions he uses to define labour do not hold up in the light of modern science. He says that the distinction between human labour and what animals do is that the humans form a plan and have a vision of what they are going to achieve *before* they do it. This sort of sharp distinction between the activity of humans and of animals, cannot be sustained in the light of modern science. It couldn't even be sustained in the light of Darwin's writings:

> *It has often been said that no animal uses any tool; but the chimpanzee in a state of nature cracks a native fruit, somewhat like a walnut, with a stone.*([14] page 28)

This tool use in chimps is subsequently well documented by Goodall[15]. It is difficult to say that animals do not consciously plan actions or do not consciously have an intention.

> *"Brehm states, on the authority of the well-known traveller Schimper, that in Abyssinia when the baboons belonging to one species (C. gelada) descend in troops from the mountains to plunder the fields, they sometimes encounter troops of another species (C. hamadryas), and then a fight ensues. The Geladas roll down great stones, which the Hamadryas try to avoid, and then both species, making a great uproar, rush furiously against each other. Brehm, when accompanying the Duke of Coburg-Gotha, aided in an attack with fire-arms on a troop of baboons in the pass of Mensa in Abyssinia. The baboons in return rolled so many stones down the mountain, some as large as a man's head, that the attackers had to beat a hasty retreat; and the pass was actually closed for a time against the caravan. It deserves notice that these baboons thus acted in concert"*.([14] page 28)

If we consider a wolf pack, dividing a task up so that some will chase the deer and others will ambush it: this is clearly intentional behaviour. They may not be able to talk, but they communicate and carry out intentional

behaviour. Hunter [16] relates how a buffalo deliberately stalks and gores a hunter. Any zoologist is bound to be able to cite many other examples.

The hard and fast distinction between humans and animals is thus untenable. What is sustainable is that humans are able to formulate and *materialise* plans for activity. There are current investigations into whether animals (at least mammalians) actually plan and use episodic memory [17]. What is incontrovertible is that humans can record and debate their plans. They can record them in language, which can be memorised using episodic memory. They can write them down using symbols and drawings. This ability affords a much more complex division of labour, which is non-instinctive in human societies.

Obviously social insects such as ants demonstrate a complex division of labour, and Marx recognises that, and says that: "*Bees put many architects to shame*". However he also says: "*Bees never form an image of what they are intending to do*". It may be true that bees do not form an image of what they are intending to do, and that they are merely following instinct, He gives another example of the spider who, he says "*puts spinners to shame*". Recent experimental work indicated that spiders are able to form plans [18, 19]. Researchers place spiders on top of pencils in a 3-dimensional space in an enclosed bottle with another insect they want on the top of another pencil. The spiders form a strategy of how to get there, and people now think that in order to build a web they are actually behaving intentionally.

It is not even feasible to say that only humans have culture, because we know that certain forms of tool use in apes [20] and macaques [21] is cultural, and we know that certain different groups of whales have local cultures, different distinct ways of hunting and co-operating. Chimpanzees adapt to local grunts when relocated [22].

None of these is the key distinction. We think the key distinctions are (1) the flexibility we get from having hands, (2) the fact that we can communicate plans to one another via language, and (3) the fact that we can read other humans' minds. There is no other species on earth, at present, which combines these traits. There are species which appear to be able to communicate (whales and porpoises) [23] but they do not have hands. There are other species with hands, but they do not have language [24]. Lieberman [25] claims that humans have a unique ability to understand the minds of others. He argues that this gives us a strong advantage over other species because it allows us to cooperate, and such cooperation leads to survival. Other animals cooperate but do not have insight the workings of the minds

of others. Other research claims that other primates have, at least in embryonic form, something analogous to the human theory of mind [26]. This idea that empathy, or a theory of mind, is an innate human trait goes back at least to Smith [27, 28]. Even if we allow for a theory of mind in other primates, the trio : *hands, empathy* and *language* are ours alone.

5. The Universal Machine

Alan Turing at first introduced the idea of the universal computer [29] as a machine that is capable of doing any mathematical operation that a human mathematician can do. It is worth looking at how he analysed what a human mathematician does, because mathematics, in his day, was taken to be the most abstract and rational type of human activity, something uniquely human at that point. He talks about someone wanting to do a calculation: they need squared paper, they write numbers down, and then they do things such as adding or multiplying them. When someone does that, they are always just looking at a small area of the sheet of paper they are concentrating on. For example, the column of digits being added up. They keep a small running total in their working memory. Every time it runs over ten, they write a little one next to the next column of digits (if you recall how you were taught to do this at school). So essentially all mathematics comes down to the mathematician being able to see a certain number of symbols at once, having a certain state of mind, and, on the basis of that state of mind, writing down new symbols. Possibly if he has an eraser, rubbing something out. So Turing says:

"Suppose I build a machine that can do the same thing?"

The first machine has a tape instead of a square sheet of paper i.e. a long strip of paper with squares on it. The machine can write numbers down on that. The machine is aware of the current number, and it is aware of its state of mind. On the basis of looking at the current number and the internal state of mind, it writes another number down, or moves the tape left or right, and changes its state of mind. So he argues: *"With this very simple machine, I capture the essence of what a mathematician is able to do, and from that I can build a universal machine."*
The universal machine is slightly different from just that particular abstract Turing machine, in that it is a Turing machine that can be given

11

instructions that enable it to emulate any other Turing Machine. This Universal Turing Machine can therefore, in principle, given the right program, solve any mathematical problem. That was what Turing's goal was when he started work on trying to develop a practical universal computer, after the war. His goal was to specify a machine that could solve any mathematical problem, including perform any proof that a mathematician could, in principle, perform. He wasn't the first person to have thought of a universal calculator. The same basic idea had occurred to Babbage in the early 19th Century [30]. He said that if he had a machine of a certain degree complexity, he suddenly realised it could solve any mathematical problem. Turing wasn't the first person to do it, but he expressed it more clearly than anyone before that. He used this idea of the universal machine to reflect back on the limits of mathematics and what mathematicians can not do. He showed that Hilbert's formalist project in mathematics, i.e. the attempt to found all of the discipline on fixed sets of axioms and rules of inference, must fail [31, 32]. Turing invented his universal computer first as a thought experiment to illuminate the limits of human mathematics. Only later does he set about building it as a practical machine [33].

It is interesting that his idea of the universal machine is published the same year that Karel Capek's play was put on the BBC. Thus the idea of a universal robot, of a universal machine, was part of the concepts of the day, and Turing is applying it in a specifically mathematical context. We know that he succeeded: he came up with designs which later generalised and we have these machines on our desks and in our pockets today. The thing about these machines is that they are not like all previous machines. Previous machines were built to achieve one given task, whereas computers can be applied to any information processing task; given sufficient time they will solve it. That has all sorts of economies, because it means only one design of computer needs to be settled on, and it can be mass produced. So, to an extent, this universality in information processing has been achieved by machinery. It must be remembered that they are disembodied intelligences, fairly limited intelligences at the moment.

At the University of Glasgow we have been working with European colleagues on the CLOPEMA project developing a cloth-folding robot [34, 35]. It is an extraordinarily difficult task for a robot even to do the things we do easily, like picking up and folding clothes. It is challenging to get a robot to even figure out the right place to pick clothes up; things that humans find trivial, even at a very young age. The automatic hand-eye coordination that
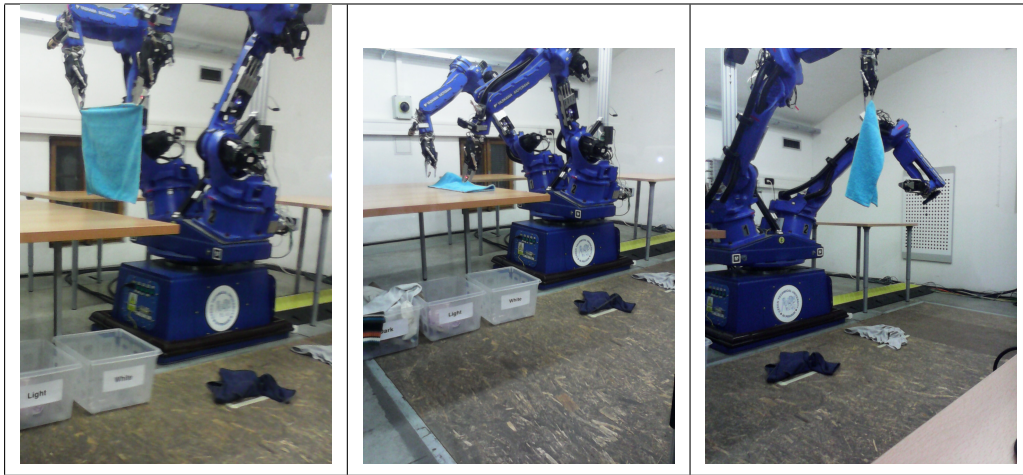
Figure 1: The CLOPEMA cloth-folding robot recognising and folding a towel. The chasis of this robot is typically a Motoman industrial unit, but it has been fitted with advanced sensing: stereo vision, touch sensitive grippers, close range cameras on its arms, force sensors in its wrists.

we have is incredibly time consuming and complicated to express mathematically. In essence we are still a long way from robots being able to do all sorts of things which are trivial human labour tasks. For example, if you have a hospital laundry, there will be specialised machines which are very good at folding clothes. They do not have general intelligence. They can fold clothes very fast, faster than a human being. They do it by laying the clothes out and having various flaps which come over to flap over the arms of a shirt. Such a special-purpose classical industrial mechanical engineering is not yet universal robotics. The engineering science of robotics is advancing at an extraordinary rate at the moment. There are robot animals, robot dogs, robot ponies, and various other animals which have been made by the robot labs at Boston, that can walk or run through the woods with the kind of gait that animals have.

Those things which have been most difficult are the things which humans find easy to do, the non-rational things. It is only very recently that robots have been capable of avoiding obstacles or moving realistically.

Our cloth-folding robot weighs a ton, and is screwed onto a concrete floor on the 7th floor of a building that we had to specially reinforce to take it, so it is a long way from something that can be deployed in just any workplace.

Just as you can use hours of labour of a standard skill level to determine the value of things, in computing, in principle, you measure the complexity of any algorithm by the number of cycles of a Turing Machine it would take. Nobody actually uses a Turing Machine, because they are relatively inefficient. It is an active research topic at the moment to see how the number of cycles of one type of machine and the length of programs you need for machines vary according to the semantic power of the machines. These things are not properly understood yet, but, in principle, we know that any computation that can be done on any other machine can be done on a Universal Turing Machine. Subject to some efficiency factor the complexity of these different calculations is equivalent. That is one of the strong hypotheses of computer science: it is not proven, but a strongly-accepted hypothesis.

6. From the laws of Thermodynamics to Values.

Essentially thermodynamics started with the attempt by Watt to improve the efficiency of steam engines. Watt already realised that heat was being converted to work, and that the more efficient your conversion of heat into work, the more efficient the steam engine. Carnot [36] subsequently shows that there is ultimately a maximum efficiency that a heat engine can achieve. In any heat engine the efficiency is related to the temperature difference between the input and output heat. Clausius [37] subsequently forms the laws of thermodynamics saying that heat always goes from a body of high temperature to one of low temperature; that you can not convert heat with complete efficiency into work, whereas Joule shows you can convert work back into heat [38]. There is an irreversibility to thermodynamic processes.

Then in the 1940s a telephone engineer working for Bell Laboratories in the States was trying to quantify the information-carrying capacity of telephone wires. Here we have something very similar to Watt. Watt was trying to quantify something which is an everyday concept — the concept of work — and he gives it a scientific meaning. Information was obviously an everyday concept prior to Shannon [39], but once you say: "How do we measure it?", you have to give it a rigorous definition, and Shannon was able to give the information content of telephone messages or telex messages a rigorous definition in terms of probability theory.

When you work through the maths of the probability theory, you come up with what is effectively the same formula as Boltzmann [40] came up with when he was trying to formulate the laws of thermodynamics in terms

of molecular motions and probability of movements of molecules. Shannon then realises that information and entropy are the same thing, and this is paradoxical because people think that information is the opposite of entropy, but what Shannon showed was that the transmission of information along a digital signal is maximized as that stream of bits becomes indistinguishable from random numbers, from random noise. The further away it is from random noise, the more redundant the encoding is, the less efficient it is at transmitting information. This represented a fundamental, huge transformation of the way people thought of information. If you read people's commentaries in the 1950s, a lot of people still did not understand it back then. Nowadays, if you are working on information theory, data compression or video compression, you treat them as the same thing. Entropy and information are essentially the same thing and that is now widely understood.

Entropy is the idea of disorder versus order. The unexpected thing is that this increase in disorder goes with increased information. If you have a hot cup of coffee in a room, for example, over time it will become the same temperature as the entire room, so it is more disordered: the heat has dissipated.

Basically, Boltzmann is saying: "What is the probability of getting something at a given state.". A state where all the heat is in the coffee cup is less probable than one in where it is being spread around the room. There are a lot more ways the thermal energy can be carried by all the gas molecules in the room, than there are ways in which it can be carried just in the coffee cup. So that, over time, the system will move to a more probable state, which is the heat spread across the room. It is from that that the maths comes out which gives you the same formula that Shannon used[4].

Shannon's idea, then, was that if you send a message down a wire, say I'm talking, and I keep on saying the words 'then' and 'and', it is not very efficient. We could express it by crunching it down into a smaller message. This was recognised even by Samuel Morse. When he devised his morse code he used a shorter series of dots and dashes for frequently-used letters than for less frequently-used letters. E is just (.) A is (. -) but Y is (- . - -). In general, if you are transmitting video data and you watch something from YouTube, that is exactly what is being done. The frequently occurring patterns in a scene can be sent with fewer bits that the more rarely occurring

---

[4]$H = \sum P \log P$

patterns. At every stage in the development of robotic perception we make use of Shannon's concepts.

To apply this to political economy, you then have to go one more step further in the development of the understanding of information that was achieved in the 60's, when Chaitin [41] was working on the problem of what we mean by a random numbers. In Russia, at the same time, Kolmogorov [42] also came up with the idea that a number is random if there is no formula shorter than that which we can generate that sequence of numbers. Chaitin later formalised this by saying: *a sequence of digits is random if there is no Turing Machine program to print it shorter than that sequence of digits which were printed out.*

This gives another definition of the information content of something. When Shannon was talking about information content and entropy, he was talking about it in terms of probability theory. Chaitin [43] gives you a different, non-probabilistic definition of information, which is based just on how few bits you need to print this information out. The shortest number of bits that you need to print something out on a computer is the information content. So, for example, he says π has a sequence of digits that go on for ever, but in fact, quite a small program, if you set it going, can go on printing the digits. Although these digits of are unpredictable unless you compute them, the actual information content is set by the length of the computer program which prints it.

This is a very deep idea related to what the nature of randomness is, and what the nature of information is. If you look at industrial production, you see that a whole series of revolutions in industrial production have occurred through information economies.

If we take one of the first ones, it is the invention of the potter's wheel. The potter's wheel enables the mass production of round pots. Why? Because a round object has low information content. All you need to do is specify the radius because that radius is shared by all points on the circumference of the pot, and the potter specifies the radius between his or her finger and thumb, as the pot is spun. You are essentially using a small amount of information to shape the whole pot.

If you do not have a potter's wheel you have to go all the way around, squishing and shaping: you do not get it perfectly round, and it takes much longer. There's a whole series of classical industrial techniques which all depend on spinning things around, so that things turned on a lathe are obviously the other main example of that. Another economy of information

which transformed things was the simultaneous application of information at right angles to the item being shaped. There are two historical industries from a long, long time ago where this type of technique was employed for mass production. The Samianware[5] industry of Gaul, in Roman Gaul, mass produced high-quality pottery with molded surfaces, so it has all sorts of Satyrs and Gods on the surface, and that is done by making a mold into which you pour the clay before it is completely hardened. Unlike the potters wheel, you are actually able to convey a lot more information to it, and you get something much more elaborate.

Instead of one point, which is the expression of the width of the pot that you do with your finger, there are many different artistic things that can be put on this one mold, and then copied.

You can use the mold repeatedly, so the human labour is initially used in carving the mold, once only. The mold subsequently transfers that information (that has taken the human a long time to produce) onto many, many pots. Another example of that very early on is the trigger mechanisms of crossbows that the Qin empire[6] used. They are mass produced, standardised trigger mechanisms. One is astonished to see samples from the Emperor's tomb, with the ceramic warriors. What is astonishing is the quality of the metalwork, the bronze swords, and the metalwork in the crossbows. These are mass produced, again using molds, to standardise size and shape.

We then come forward to Babbage. After Adam Smith, Babbage, was the next economist, who was really concerned about the basis of industrial productivity. Smith had been concerned about it and explained it terms of the division of labour. Babbage actually had superior engineering knowledge to Smith, and in his book [30] he identifies other principles than the simple division of labour which lead to increases in production, and the techniques of copying he identifies as being absolutely key.

Another industry which produced a huge transformation was the printing industry. If we analyse that in terms of information theory and entropy, what you start off with is just random cellulose fibers which are in a mush and they are completely disordered. You boil it up, and then you roll it out to

---

[5] http://www.britishmuseum.org/explore/highlights/highlight_objects/pe_prb/s/samian_ware_pottery.aspx

[6] https://www.boundless.com/art-history/textbooks/boundless-art-history-textbook/chinese-and-korean-art-before-1279-14/the-qin-dynasty-96/art-of-the-qin-dynasty-459-5603/

form paper, and this process of rolling it out is a process of reducing the entropy of the material: it has few degrees of freedom. Fibres in paper are forced to align with the surface, and therefore have fewer degrees of freedom.

There is less randomness in it, it is more ordered; it has lower entropy. Having produced a low entropy material, you can then add human information to it. You've reduced the natural information in it, and you replace it with human information when you bring down the printing head and write letters on it. This printing head coming down is the same technology that is used in molding. The printer head comes down and it puts a dye onto the paper and now the paper is more disordered between white and black. Essentially we've got more disorder in that paper.

Another anecdote serves to illustrate this point. Borges writes about the universal library [44] where every possible book exists, with is every possible permutation of letters. Obviously the great bulk of those means nothing, but some do mean something. There is a huge number of possible permutations of letters that a printer can put on a paper, but he puts one particular combination, which is a human-specified combination.

If you look at a lot of manufacturing processes, they start out by reducing the entropy of a material, and then they add some kind of human-specified information to it. It is the same process that takes place to make a car. A lot of energy has to be used to reduce the entropy of iron-ore to make iron, and more work is then done to reduce the entropy of a large block of iron and turn it into a flat, rolled sheet. Given sheet-steel, which can be pressed into any shape, a particular die will force it into a car door shape. It is doing the same thing as the Samianware does: conveying information in parallel. If we look at the mass production industries, they are dominated by industries which are able to transfer human-generated information, in parallel, at right angles onto the substance. The industry which has had the hugest improvement, the semiconductor industry, is essentially a printing industry — a micro printing industry.

Essentially all this becomes evident once you digitise production. Suppose you are printing something, you actually know how many bits you have to send to the laser-printer to print the information. Take the different types of laser-printers. With the first generation of laser-printers you had to send every bit to the laser-printer at 300 dots per inch: you had to switch the print head on and off, on and off, on and off. Then Adobe invented postscript. What postscript does is that it essentially turns your laser into a Turing Machine. What you send to the printer is a postscript program that then

prints your image. Every time you use an Adobe postscript laser-printer, you are applying Chaitin's theorem. You are sending less information than is naïvely required. because you are sending a program to print the whole thing.

7. Evidence that Labour is still Important

There is a lot of empirical evidence that supports the classical *Labour Theory of Value*. Orthodox economics pays no attention to it, doesn't believe it is true at all, and doesn't take it seriously. But it is actually very well supported empirically. Consider almost any country. If you compute the amount of labour required to produce the output of every industry, you can see by how much the output of that industry sells. There will be more than a 95% correlation coefficient between the two. This means that 95% of the variation in prices of the value of the output of the industries is determined by the amount of labour required to produce that output [45, 46, 47, 48].

It is an extraordinarily strong result. There are not many pieces of economics which are as well attested. The classical economists like Smith and Ricardo took it as so obvious that they didn't need to produce the evidence. When neo-classical economics replaced classical economics as the dominant theory, they didn't actually produce any empirical evidence against the Labour Theory of Value: they constructed an abstract, non-empirical set of mathematical theories, which they claimed would explain how prices operated. Unfortunately these are actually a non-testable set of theories because they have got more free variables than things you can observe. If you've got a theory with more free variables than observations, then, by twiddling the free variables you can explain anything, and your theory explains nothing because it explains everything. This is Chaitin's point again: scientific theories have to have an information economy; they have to predict more than the content of the theory [49].

This is the idea of Occam's Razor — the shorter the explanation, the more concise it is, the more, in some sense, correct that theory is. Basically, neo-classical theory has no predictions of the overall price structure of the economy. So why does the *Labour Theory of Value* still hold?

You can empirically observe that it holds and you can make hand waving analogies between human labour and energy as Marx did, but that doesn't explain why it is that labour determines the value of things.

One clue is to say that human beings are the only *universal* resource we have, this would be a Capek/Turing style argument. Another theory was put forward by two mathematicians Farjoun and Machover[7] [50] in the 1980s.

Essentially they used a statistical mechanics argument to show that if you make certain quite plausible assumptions about the probability of a firm surviving if the price it obtains for its product is less than what they called the integrated labour coefficient of that product. Suppose you are making chocolate digestive biscuits, and a certain amount of labour is needed to make the chocolate digestive, and wages have to be paid for that. Wages have to be paid for the labour that goes into supplying the chocolate, that goes into supplying the flour etc. You can go back and see how much was paid for labour all the way back to the farmer. They say: "Suppose a firm is only able to sell chocolate digestive biscuits for less than all the wages that were paid all the way back through the supply-chain'. Well, It is clear that it is very unlikely that the firm will remain in business for long, because not only is it making a loss on its own account, but it is not even able to meet the wage bills of the previous layers of the manufacturing process. So they said: "Let's assume only $\tilde{5}\%$ of firms can be operating in that position. Then let's make a parsimonious assumption and assume that the ratio of actual prices to labour content is normally distributed."

If it is normally distributed then only a certain percentage of firms can be in the position where they are selling it at such a low price. Then you can say we have constrained the normal distribution to have a certain standard deviation. On that basis, they make a prediction about the standard deviation of the ratios of values to prices, and they say: "*It must be narrowly distributed*". Any system with a large number of free variables which are added together to produce an effect ends up being normally distributed: that's just a property of probabilities. In essence a random system, an entropic system, will have a normal distribution. Now, if we just assume that, and we make the additional constraint that firms can not remain solvent if they can not meet their costs, with a low probability of surviving, then that is enough to generate the labour theory of value.

The argument is basically an entropic one. The selling of commodities is a chaotic process which ends up, due to entropy maximisation, with there being

---

[7]Interestingly Machover is a recursion theorist, i.e. works on the same general branch of maths as Turing did.
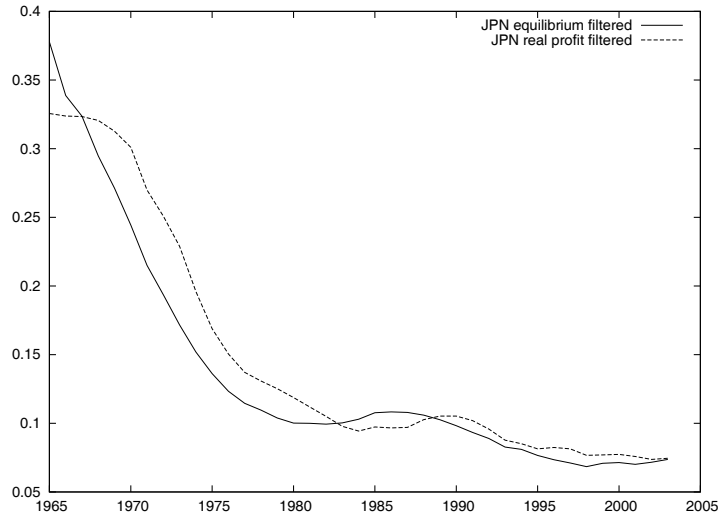
Figure 2: The rate of return on capital in Japan where robots are widely used has fallen very low. The dashed line is the empirically measured rate of return, and the solid line the equilibrium rate of return predicted by the labour theory of value, that is to say the center of gravitation for the empirical rate. Graph by Tamerlan Tadjadinov, `http://compbio.dcs.gla.ac.uk/cgi-bin/profits/home.cgi`

a particular probability distribution of prices to labour content. This entropy maximising distribution is one which has prices relatively closely clustered around labour values. Watt and Smith were concerned with work, Watt and Marx with work and power. Concepts derived from thermodynamics go from explaining steam engines with Clausius, explaining gases with Boltzman, explaining information with Shannon to deducing the labour theory of value with Farjoun and Machover [51]. We come around full circle.

8. Second Cycle of Mechanisation

The whole foundation of the validity of labour theory of value has been the indispensability of human labour. What will happen when robots can replace almost all human workers?

If there are going to be more robots than people, it would seem that the capital labour ratio in the economy will go through the roof. Both Marxian

21

[52, 53] and neo-classical theory [54] predict that under such circumstances the rate of return on capital, the rate of profit, will be much lower. Indeed if we look at Japan where robots are perhaps more widely used than any other country, the rate of return on capital has fallen very low (Figure 2). This would seem to bear out the prediction.

A low rate of profit has a depressing effect on a capitalist economy. It discourages investment and indeed the Japanese economy has been mired in stagnation since the 1990s. There is a fear that this process is spreading across the developed world [55]. Summers, in an influential article, talked about a fall in the 'natural rate of interest' [56] contributing to 'secular stagnation', and Roberts [57] refers to a fall in the rate of profit leading to the long recession since 2008.

Robert Reich, former US Secretary of Labor, describes the development of a nightmare economy in which robots do all the predictable work and humans are limited to the scraps of work that robots cannot yet do [58]. These unpredictable micro-jobs will be coordinated by software packages such as those used by Amazon for their 'Mechanical Turk' service or the Über taxi hire system. It is an economy with no job security, no workplace pension or sickness benefit. The consequence is growing inequality, a middle class driven into debt and unable to provide the market for the growing mass of products produced by robot industry, which all reinforces secular stagnation [59]. This narrative has been greatly strengthened by the recent publication Piketty's monumental work [60] on growing inequality.

There are some problems with the narrative. For a start, Piketty's model for the growth of inequality is based on the return on capital being greater than the rate of growth of the economy. Under these circumstances he argues the share of wealth going to the owners of capital is bound to increase from decade to decade, leading to a more and more unequal society. If Japan is anything to go by, intense automation will lead to a lower rate of return and tend to stabilise the degree of inequality of *wealth*. It remains plausible, though, that automation weakens the bargaining position of labour and results in a more uneven distribution of *income*.

Even in Japan the robots in use today are not general purpose. Human labour is still one essential input into all production processes. Japan has an aging population and a low birth rate, which provides an incentive to use robots, but these robots are nothing like the pass-as-human Androids in Blade Runner. The look much like the one shown in Fig 1. They can not walk off the production line and take a job driving a bus or working in a

hospital. The whole argument for robots leading to a high capital labour ratio and thus a low rate of return is the presumption that human labour remains an essential input to every industry.

Suppose real droids of the sort envisaged by Asimov, Capek etc. could be built: would this still hold? Could you, at least in principle, still have a viable capitalist economy? For an answer to this we can turn not to Turing who famously argued that computers will become as intelligent as humans [61], but to another great pioneer of computing, John Von Neumann. Turing was a polymath genius but his interests did not extend to economics. Also a polymath, Von Neumann's contributions spanned quantum mechanics [62], computer design [63], self replicating machines [64] and economics [65]. He came up with a basic design of computer still copied in every laptop and smartphone.

What concerns us here are two other contributions. His theory of economic equilibrium [65] described an idealised expanding capitalist system that produced outputs in the same proportions as the inputs it used up. He showed that under these circumstances the rate of growth would be the same as the rate of profit and that there would be a unique set of prices that would allow all industries to make this same rate of profit. His model economy was represented by what later came to be called an input output table and which can be considered a generalisation to many more sectors of the two sector *reproduction schemes* invented by Marx [66]. In Von Neumann's model, labour is just another produced input — produced by the amount of food, clothes etc. required for the worker's survival: a very classical or Ricardian conception. It is clear that, in principle, without doing violence to his Maths one can use the same sort of equations to represent an economy in which all the 'labour' is now provided by robots. What you would then have is the *reductio ad absurdam* of capitalism. It would be an economy of competing firms, staffed by robots, all of whose output went into producing more machinery and more robots. The robot-operated firms could still buy and sell things, there would be prices and profits but there would be no human consumers. It would be a capitalism unconstrained by the labour shortages that have slowed Japan down, expanding until it met the physical limits of the Earth.

Another strand of his work, addressed the problem of what would be required to have systems of self-replicating machines [64], machines able to build themselves without human intervention. This spawned off the field now known as cellular automata which underly computer games like "Life" [67] or

the early releases of of the popular economic simulation game *Sim City*[68]. More significantly his model of self reproducing machines which contain software instructing them to make more copies of themselves was hugely influential in the modern understanding of the living cell and the role of the genetic code in the cell. From that it goes on to influence the philosophy of evolutionary theory in Dawkins [69] or Dennett [70].

So, in principle, one could have an exponentially growing robot capitalism. Obviously a system that devotes all its output to producing more robots is not in human interest. However, it would be a simple matter to alter the Von Neumann economic model to get one that described a non-growing rentier capitalism, one in which the entire surplus of the robot operated industry was directed at producing luxuries for their owners rather than being reinvested[8]. At this point, where universal robots are available, human labour would have ceased to play the role that Adam Smith identified for it of regulating all prices. Profit could be made without them. Workers would become functionally redundant to the corporate economy. They would be the Nazis' useless mouths. European history shows what happens when a state treats humans as dispensable.

9. Reflection

When Stephen Hawking told the BBC: "*The development of full artificial intelligence could spell the end of the human race*". he was apparently thinking in terms of the intelligent machines themselves being our enemy. It should be borne in mind that machines exist in the context of systems of legal and social relations. We conventionally think of these relations as being relations between people, but 'people' has an elastic meaning in this case. Companies are legal persons, juridical subjects and states subjects in international law. The Pashukanis school of legal theory [72] argued that it is the abstract structural relations that come first. Legal personalities grow from the relations. Humans are the bearers of these relations: actors playing roles in a script laid down by the unwritten laws of social relations. In a metaphor from Greek drama, Marx describes capitalists as *Träger* or wearers of the capitalist mask, as actors bore masks in classical drama. Such a

---

[8]We feel dystopian models like that of Benzell *et al.* [71], which ignore the existence of firms owning robots and assume all investment comes out of worker's savings, are too far removed from capitalist reality to provide a guide to what could happen.

mask can, in principle, be worn by a computer. Indeed we already see this when algorithmic traders enter into contracts that they have calculated will be profitable. In financial markets time critical decisions are already being taken by machine. As long ago as the 1987 algorithmic trading was being blamed for a major stock market crash [73]. Given that such crashes affect the lives of millions[9] one could argue that the robots already took over 25 years ago.

From Turing's standpoint, a network of human traders applying profit maximising rules, or a network of computers performing algorithmic trades are computationally the same, differing only in their clock speed. Indeed there is little statistical evidence that the 1987 crash was much different, other than in speed, from earlier ones when trading was manual [74]. Ever since Hyndman blamed the severity of a crash on the telegraph [75], there has been a tendancy to blame the latest communication technology. The traders, whether manual or automatic, collectively comprise a distributed computer with complex, irreducible, emergent behaviour.

The metaphor of the economy as a vast machine, crushing humans, indifferent to lives, hopes and aspirations was well understood by artists in the 20s and 30s: Capek's *RUR*, Lang's *Metropolis*, Chaplin's *Modern Times* all expressed this theme. The question is: how can the little people, Chaplin's heroes, overcome it?

10. Conclusion

The popular press trumpets the fact that robots will take over the world [76, 77]. This paper considers the viability of this prediction by constructing a philosophical argument, mining the rich literature on economics, information theory and thermodynamics. We conclude that robots do not constitute a serious threat, in terms of replacing humans in the labour market until such time as genuinely 'universal' robots are made.

[1] P. K. Dick, Do Androids Dream of Electric Sheep, Grenada, 1968.

[2] I. Asimov, I, Robot, Gnome Press, 1950.

[3] K. Vonnegut, Player Piano, Scribner, 1952.

---

[9]One of the authors lost their job within days as a result of the crash.

[4] R. Cellan-Jones, Stephen Hawking warns artificial intelligence could end mankind, BBC.
URL http://www.bbc.co.uk/news/technology-30290540

[5] K. Marx, Capital, Vol. 1, Progress Publishers, Moscow, 1954, original English edition published in 1887.

[6] D. Ricardo, Principles of political economy and taxation, G. Bell and sons, 1891.

[7] M. Brooksbank, Sidlaw Breezes, Unknown, 1952.

[8] D. S. L. Cardwell, From Watt to Clausius: the rise of thermodynamics in the early industrial age, Cornell University Press, 1971.

[9] A. Smith, The Wealth of Nations, Penguin, 1974.

[10] P. Cockshott, A. Cottrell, Labour time versus alternative value bases: a research note, Cambridge Journal of Economics 21 (1997) 545–549.

[11] P. Sraffa, Production of commodities by means of commodities, Cambridge University Press, Cambridge, 1960.

[12] K. Capek, Rossum's Universal Robots, Dover Publications, 1999.

[13] C. S. C. Lim, Designing inclusive ICT products for older users: taking into account the technology generation effect, Journal of Engineering Design 21 (2-3) (2010) 189–206.

[14] C. Darwin, The descent of man, Digireads.com Publishing, 2004.

[15] J. Goodall, Tool-using and aimed throwing in a community of free-living chimpanzees., Nature 201 (1964) 1264.

[16] J. A. Hunter, Hunter, Safari Press, 2000.

[17] N. S. Clayton, T. J. Bussey, A. Dickinson, Can animals recall the past and plan for the future?, Nature Reviews Neuroscience 4 (8) (2003) 685–691.

[18] D. Harland, D. Li, R. Jackson, How jumping spiders see the world, How Animals See the World: Comparative Behavior, Biology, and Evolution of Vision (2012) 133–164.

[19] J. McCrone, Smarter than the average bug, New Scientist 190 (2553) (2006) 37–39.

[20] A. Whiten, J. Goodall, W. C. McGrew, T. Nishida, V. Reynolds, Y. Sugiyama, C. E. Tutin, R. W. Wrangham, C. Boesch, Cultures in chimpanzees, Nature 399 (6737) (1999) 682–685.

[21] S. Kawamura, The process of sub-culture propagation among japanese macaques, Primates 2 (1) (1959) 43–60.

[22] J. Webb, Chimps 'learn local grunts' to talk to new neighbours, http://www.bbc.co.uk/news/science-environment-31128237 (February 2015).

[23] K. McComb, D. Reby, L. Baker, C. Moss, S. Sayialel, Long-distance communication of acoustic cues to social identity in African elephants, Animal Behaviour 65 (2) (2003) 317–329.

[24] R. Dunbar, Why only humans have language, Lucy to Language: The Benchmark Papers (2014) 427.

[25] M. D. Lieberman, Social cognitive neuroscience: a review of core processes, Annu. Rev. Psychol. 58 (2007) 259–289.

[26] C. M. Heyes, Theory of mind in nonhuman primates, Behavioral and Brain Sciences 21 (01) (1998) 101–114.

[27] A. Smith, The theory of moral sentiments, Penguin, 2010.

[28] D. Göçmen, The Adam Smith Problem: Reconciling Human Nature and Society in 'The Theory of Moral Sentiments' and 'Wealth of Nations', Vol. 1, Tauris Academic Studies, 2007.

[29] A. Turing, On Computable Numbers, With an Application to the Entscheidungsproblem, Proceedings of the London Mathematical Society 42 (1937) 230–65.

[30] C. Babbage, The Economy of Machinery and Manufactures, Charles Knight, Pall Mall East, London, 1832.

[31] G. J. Chaitin, The unknowable, Springer Science & Business Media, 1999.

[32] P. Cockshott, L. M. Mackenzie, G. Michaelson, Computation and its Limits, Oxford University Press, 2012.

[33] J. Copeland, Alan Turing's Electronic Brain, Oxford, 2005.

[34] J. Stria, D. Prusa, V. Hlavác, L. Wagner, V. Petrik, P. Krsek, V. Smutny, Garment perception and its folding using a dual-arm robot, in: Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on, IEEE, 2014, pp. 61–67.

[35] L. Sun, G. Aragon-Camarasa, P. Cockshott, S. Rogers, J. P. Siebert, A heuristic-based approach for flattening wrinkled clothes, in: Towards Autonomous Robotic Systems, Springer, 2014, pp. 148–160.

[36] S. Carnot, Reflections on the Motive Power of Fire, Robert Fox, 1824.

[37] R. Clausius, On the moving force of heat, and the law regarding the nature of heat itself which are deducible therefrom, Annalen der Physik 2 (1-21) (1851) 102–119.

[38] J. Joule, On the mechanical equivalent of heat.[abstract], in: Abstracts of the Papers Communicated to the Royal Society of London, JSTOR, 1843, pp. 839–839.

[39] C. Shannon, A mathematical theory of communication, The Bell System Technical Journal 27 (1948) 379–423 and 623–56.

[40] L. Boltzmann, Lectures on gas theory, Dover Pubns, 1995.

[41] G. Chaitin, Information and randomness: A survey of algorithmic information theory, in: The Unknowable, Springer, Signapore, 1999.

[42] A. Kolmogorov, On tables of random numbers, Theoretical Computer Science 207 (2) (1998) 387–395.

[43] G. J. Chaitin, On the length of programs for computing finite binary sequences, Journal of the ACM 13 (4) (1966) 547–569.

[44] J. L. Borges, The total library, in: E. Weinberger (Ed.), Selected nonfictions, Viking, 1999, pp. 214–6.

[45] A. M. Shaikh, The empirical strength of the labour theory of value, in: R. Bellofiore (Ed.), Marxian Economics: A Reappraisal, Vol. 2, Macmillan, 1998, pp. 225–251.

[46] D. Zachariah, Labour value and equalisation of profit rates, Indian Development Review 4 (1) (2006) 1–21.

[47] P. Petrovic, The deviation of production prices from labour values: some methodolog and empirical evidence, Cambridge Journal of Economics 11 (1987) 197–210.

[48] E. M. Ochoa, Values, prices, and wage–profit curves in the us economy, Cambridge Journal of Economics 13 (1989) 413–29.

[49] P. Cockshott, A. Cottrell, The Scientific Status of the Labour Theory of Value, IWGVT conference at the Eastern Economic Association meeting, in April.
URL http://www.helmutdunkhase.de/eea97.pdf

[50] E. Farjoun, M. Machover, Laws of Chaos, a Probabilistic Approach to Political Economy, Verso, London, 1983.

[51] E. Farjoun, Production of commodities by means of what?, in: E. Mandel (Ed.), Production of commodities by means of commodities, Cambridge University Press, 1984.

[52] K. Marx, Capital: Critique of Political Economy, Vol. 3, The Process of Capitalist Production as a Whole,, The Marx/Engels Internet Archive, 1894.

[53] A. Kliman, Reclaiming Marx's Capital: a refutation of the myth of inconsistency, Lexington Books, 2007.

[54] R. M. Solow, Capital theory and the rate of return, North-Holland Amsterdam, 1963.

[55] A. S. Posen, The realities and relevance of japan's great recession: neither ran nor rashomon, Peterson Institute for International Economics Working Paper (10-7).

[56] L. H. Summers, Us economic prospects: Secular stagnation, hysteresis, and the zero lower bound, Business Economics 49 (2) (2014) 65–73.

[57] M. Roberts, From global slump to long depression, International Socialism 140 (2013) 17–42.

[58] R. Reich, The share the scraps economy (February 2015).
URL robertreich.org/post/109894095095

[59] B. Z. Cynamon, S. M. Fazzari, Inequality, the great recession, and slow recovery, Available at SSRN 2205524.

[60] T. Piketty, A. Goldhammer, Capital in the twenty-first century, Belknap Press, 2014.

[61] A. Turing, Computing machinery and intelligence, Mind LIX (236).

[62] J. Von Neumann, Mathematical Foundations of Quantum Mechanics, Engl. transl. of the 1931 German edition by RT Beyer (1955).

[63] J. Von Neumann, First Draft of a Report on the EDVAC, IEEE Annals of the History of Computing 15 (4) (1993) 27–75.

[64] J. Von Neumann, A. Burks, et al., Theory of self-reproducing automata, Univ. of Illinois Press Urbana, IL, 1966.

[65] J. Von Neumann, A Model of General Economic Equilibrium, Review of Economic Studies 13 (33) (1945) 1–9.

[66] K. Marx, Capital, Vol. 1. The process of production of capital, Trans. S. Moore and E. Aveling, Ed. F. Engels. Moscow: Progress Publishers. URL (accessed December 2007): Marx/Engels Internet Archive http://www. marxists. org/archive/marx/works/1867-c1, 1887.

[67] J. Conway, The game of life, Scientific American 223 (4) (1970) 4.

[68] J. A. Betz, Computer games: Increase learning in an interactive multidisciplinary environment, Journal of Educational Technology Systems 24 (2) (1995) 195–205.

[69] R. Dawkins, The selfish gene, Oxford university press, 2006.

[70] D. C. Dennett, Darwin's dangerous idea, The Sciences 35 (3) (1995) 34–40.

[71] S. G. Benzell, L. J. Kotlikoff, G. LaGarda, J. D. Sachs, Robots Are Us: Some Economics of Human Replacement, National Bureau of Economic Research, 2014.

[72] E. B. Pashukanis, Law and Marxism: A General Theory Towards a Critique of the Fundamenta Juridical Concepts, Pluto Publishing, Worcester, 1989.

[73] J. P. Raines, C. G. Leathers, The new speculative stock market: why the weak immunizing effect of the 1987 crash?, Journal of Economic Issues (1994) 733–753.

[74] G. W. Schwert, Stock market volatility: Ten years after the crash, Tech. rep., National Bureau of Economic Research (1998).

[75] H. M. Hyndman, Commercial crises of the nineteenth century, Routledge, 2013.

[76] E. ZOLFAGHARIFARD, Is the world heading towards a post-human future? sir martin rees warns that super-intelligent robots could wipe out humanity, daily Mail (21 May 2015).

[77] M. Rees, Astronomer royal martin rees: How soon will robots take over the world?, the Telegraph (23 May 2015).