BOETTKE, SYNTAX AND THE TURING TEST

1. Boettke's Argument

Philosophy, wrote Althusser,(Althusser 1971b) represents politics to science and science to politics. It draws demarcation lines around permited knowledge and permited investigation. This is particularly true of philosophical interventions by economists like Hayek and Boettke(Boettke and Subrick 2001). Hayek's overarching project was to provide a justification for his claim that we face insurmountable limits in our ability to control complex phenomena - in particular the economy. Boettke wants to use the philosophical arguments of Searle against AI to rule out as scientifically impossible any proposal to establish a modern planned economy utilising computer technology on the basis of a philosophical argument about the nature of knowledge and the limits of computation. His argument is allusive rather than rigourous, but even within these limits we think that it rests upons some fundamental mis-understandings relating to the nature of information and of computation.

Boettke deploys Searle's critique of strong versions of artificial intelligence to ampliphy Hayek's critique of socialist planning, to clarify the deeply problematic transition from planned economies in Eastern Europe and the former USSR. Boettke's argument may be summarised as follows:

- Searle argues that there is a fundamental distinction between syntax and semantics.
- Hayek argues that there is fundamental distinction between formalisations of economic organisations and the underlying social bases of real economies.
- Thus changing the 'syntax' of the economy from planned to market led without changing the 'semantics' of the underlying social formation has resulted in the chaotic form that capitalism took in these societies during the period of transition.

Boettke deploys this argument in the context of a critique of Cockshott and Cottrells proposals (Cottrell and Cockshott 1992, Cockshott and Cottrell 1997)for a revived form of socialist economy based on distributed planning. Unlike other critics of planning(Nove 1983) Boettke conceeds that the computational problems of planning are no longer intractable. Rather he argues, like Hayek, that it is impossible to fully capture the 'meanings' necessary for a planned economy in a set of rules independent of the human actors who constitute the economy.

We are materialists in a Marxist tradition. We are also materialists in the Turing tradition. We will deploy arguments from both these traditions in uur critique of Boettke. Our arguments in the following sections may be summarised as follows:

- We dispute the fundamental distinction drawn by Searle between syntax and semantics.
- We would also dispute the unproblematised transfer of concepts like syntax and semantics from the domain of formal language theory to the quite different domain of socio-economic phenomena.

- As against Hayeks's emphasis on the centrality of human subjectivity and the limits that this places on social science, we counterpose a model in which human subjects are actors scripted by the social relations of production.

2. Formal semantics and inetrpretation

Boettke's account of Hayek, Searle and the social context of economic activity uses "semantics" in three mutually contradictory senses, but let us return to that later. His main argument centres on the assertion that symbols are not adequate for capturing semantics. In response, we will draw on the remarkable contributions of the British mathematician Allan Turing, whose pioneering work encompasses both the formal characterisation of semantics and the continuity between artificial and natural intelligence. To do so, we will need to briefly recapitulate paradigmatic moments in mathematical logic and philosophy, in the 1930's and 1950's respectively.

The German mathematician Hilbert enunciated an ambitious Programme of putting all of mathematics on a formal basis, in particular showing that it was consistent (incapable of proving contradictions), complete (capable of establishing all true theorems) and decidable (encompassing a mechanical means of establishing whether or not an arbitrary formula was true or false). Russell's paradox showed that an adequately expressive formalism for mathematics based on set theory could certainly express contradictions. The only escape was through a seemingly infinite regress of more and more powerful systems. Godel's theorems showed that an adequately expressive formalism for mathematics could be either consistent or complete, but not both. In seeking consistency, mathematicians had to accept that there were theorems whose truth could not be established. Incidentally, Boettke refers to Penrose's illegitimate use of Godel's results in trying to establish that human beings are more than machines., which we discuss below.

Both Russell's paradox and Godel's theorems depend on the formal language of mathematics being strong enough to describe itself. As we shall see, this property is key to our refutation of Boettke's claim that semantics is beyond symbols.

Turing approached the problem of decidability from the angle of computability, that is of seeking a characterisation of those mathematical results which could be established by mechanically applying rules to symbol sequences. He devised an astonishingly simple rule following machine, now known as a Turing Machine (TM) (Turing 1937).

A TM has an infinite tape composed of a sequence of cells. Each cell can hold an arbitrary symbol. In practice, "0" and "1" suffice. There is a notional reading head which can inspect one cell on the tape. After inspection, the tape may be moved one cell to the left or right under the head. A TM is controlled by transition rules. At any given moment a TM is in a definite state, that is there is a particular subset of the transition rules which it is appropriate to consider considered. Each rule says:

if the TM is in state X with symbol Y on the cell under the reading head then it should replace the symbol Y with a new symbol Y', move the tape one cell either left or right, and change to state X'.

Now, an arbitrary set of rules may itself be described by a set of symbols on a TM tape. Turing used this property to devise a Universal TM (UTM), that is a TM that could simulate an arbitrary TM from the latter's description on a tape. This result should not seem so strange to those familiar with DNA's ability to encode itself. Here we have a fundamental biological property reflected in a abstract mathematical artefact.

To return to Hilbert's Programme, Turing used his TM to establish the limitations of formal characterisations of computability. He showed that, even though it is possible to construct a UTM that will behave like an arbitrary TM, it is not possible to construct a TM which can tell if an arbitrary TM will ever halt on an arbitrary tape. Once again, this disconcerting result depends on the ability of a formal system to describe itself.

Many other people studied computability, and established their own formal models quite unlike TMs, for example Kleene's recursive function theory, Church's lambda calculus and Markov's transition systems. Curiously, all these characterisations are equivalent to TMs and to hence to each other, in the sense that any computation that can be described in one formalism can be described with equal accuracy in any other formalism. Such characterisations are said to be Turing Complete (TC). The Church-Turing thesis is that all future accounts of computability will be TC. So far this has indeed been the case. In particular, TMs are equivalent to von Neumann machines with infinite memory, that is common or garden digital computers.

It is important to note that Turing's motivation was also to characterise the limitations of humans performing computations. In particular, he wished to show that there are computations which humans cannot carry out and the TM was intended to correspond directly to a human performing a computation. Suppose a person has a pencil and a sheet of paper,is familiar with school arithmetic and wishes to multiply two 8-digit numbers. At any given moment they can only remember a small part of the computation in their brains and have to use the paper to record intermediate results. They will know their times tables and be able multiply a given pair of digits from memory to give a one or two digit result. They write down the two numbers, one above the other, and then start by multiplying the top number by the least significant digit of the bottom number. In turn, to do this they start by multiplying the least significant digit of the top number by the least significant digit of the bottom number to give a two digit answer. They write down the least significant digit of the answer, remember the carry digit, multiply the next most significant digit in the top number by the least significant digit in the bottom number, add in the carry from the previous stage and so on. Thus, at any stage they only have to remember one or two digits, which correspond to the distinct internal states of a TM. The sheet of paper corresponds to the TM tape and the times tables correspond to the TM state transition rules.

A mathematician proving a theorem operates under similar constraints. They can only remember a small psrt of the proof at any stage and have to use pencil and paper to record the intermediate steps. Furthermore, at any stage they can only choose from a fixed set of rules of inference in carrying the proof forwards.

Let us now turn to semantics. Modern accounts of formal semantics are based on Tarski, who uses a meta-language, a language for describing other languages, to link general syntactic constructs in the language being defined to their meanings. Tarski's conception of semantics is declarative; that is symbol sequences ultimately denote values in some abstract space. In contrast, Wittgenstein gives a more procedural account of semantics as involving the application of rules to symbol sequences.

He places particular emphasis on the role of the interpretation of rules in language games, where communicating enitities must carry out consistent interpretations of shared rules on symbol sequences in order to share the same meanings. Thus there are three levels to semantics: the symbol sequences whose meaning is to be established, the rules for establishing meanings and the mechanism for applying rules to symbol sequences.

How then are mechanisms for applying rules to be characterised? We seem to enter an infinite regress of layers of meta-rules. However, Turing's UTM shows that meta-rules are no different from any other rules and may in turn be encoded as symbol sequences. Thus, there is no need for teleological recourse to semantics beyond symbols, provided there are *material devices* which can interpret symbol sequences as self encoding rules over other symbol sequences. Digital computers are one such class of devices. Cells full of DNA are another. Human brains are a third.

Boettke lays great emphasis on the importance of context in understanding. This is then treated as analogous to the distinction between syntax and semantics, with the claim being that computers operate purely syntactically and thus neglect context. But this is to use too simple a notion of syntax. A syntax that ignores context is certainly very limited. Chomsky (Chomsky 1956) distinguished 3 different models of grammar that a language might have, ranging from the simplest Finite State Grammar, through Phrase Structured Grammars to the most complex: Transformational Grammar. He argues that only the last is poweful enough to represent the English language.

Corresponding to each of these levels of grammar there is a class of automata that can act as their recogrnisers(**?**). Finite State automata can recognise Finite State Grammars and push-down automata can recognise Phrase Structured Grammars. For example, simple four function pocket calculator is based on a finite automaton and the commands that you can give it on the keypad constitute a Finite State Language. A more sophisticated scientific calculator that supports bracketed formulae, uses a push-down automaton. Chomsky's Finite State and Phrase Structured Grammars are *context free* systems for the generation or derivation of sentences, and the languages they produce are termed context free languages.

Boettke says :

> Rules are not enough. Symbols are not enough, What is needed is an understanding of how the rules interact and evolve with the symbols over time.

As a characterisation of the limits of context free languages this is fair enough. But the languages, like C++ or Java used to program general purpose computers (UTM equivalent ones) do not use context free grammars. In these languages there exist phrases whose meaning can only be decided if one takes into account the context within which the phrase is embedded. They allow you introduce new words and whole new classed of entities into a text. The meaning of these words can only be determined by examination of remote parts of the text, or indeed by reference to quite other texts in which the words or entities were defined.

If Boettke were correct, computers would be incapable of handling these context sensitive languages. But they can. Provided that the computer has access either to a random access memory or the sort of read/write tape proposed by Turing context sensitive languages can be reliably parsed.

Boettke's claims about the contextual limits of computers are only valid when applied to machines below the complexity of UTMs.

3. Intelligence and the Chinese Room

Turing's second major contribution was to Artificial Intellegence(Turing 1950), in framing the "Turing Test" for determining whether or not a behaving entity is intelligent.. Boettke draws on Searle's refutation of Turing through the "Chinese Room" argument. Again, before discussing Searle we will summarise Turing's Test. A good overview of the literature on this is provided by Hauser[1].

Turing starts by discussing the "Imitation Game" where an interogator has to decide which of two people is male and which is female simply by asking them questions. To make this harder, the interrogator cannot see or hear the people but interacts with them by passing notes [2].

Turing then considers playing the game with an entity which may be human or a machine where communication is by teletype. His "Test" succeeds if the machine succesfully passes itself off as human. The Turing Test is important for proposing a purely behavioural measure of intelligence without reference to internal mechanism. Turing points out that this is precisely why we attribute intelligence to other people; it being impossible to in any sense know what it its like to be another person. In denying a behavioural basis for attributing intelliegence, we are reduced to solipsism.

Turing argues strongly that there is no principled reason why it should not be possible to construct a machine that passes the Test. He presents and refutes seven arguments against Artificial Intelligence, in particular the Mathematical Argument. This claims that undecidabilty results show that there are limitations to formal computations which are not shared by human beings, who are able to recognise such limitations. Boettke alludes to Penrose's development of this argument. Turing's response is:

> ... although it is established that there are limitations to the powers of any particular machine, it has only been stated, without any sort of proof, that no such limitations apply to the human intellect. But I do not think this view can be dismissed quite so lightly. Whenever one of these machines is asked the appropriate critical question, and gives a definite answer, we know that this answer must be wrong, and this gives us a certain feeling of superiority. Is this feeling illusory? It is no doubt quite genuine, but I do not think too much importance should be attached to it. We too often give wrong answers to questions ourselves to be justified in being very pleased at such evidence of fallibility on the part of the machines. Further, our superiority can only be felt on such an occasion in relation to the one machine over which we have scored our petty triumph. There would be no question of triumphing simultaneously over all machines. In short, then, there might be men cleverer than

[1] in Searle's Chinese Box: Debunking the Chinese Room Experiment, Minds and Machines 7:199-226, 1997 (http://members.aol.com/lshauser2/chinabox.html).

[2] This may have been particularly pertinent to Turing, a homosexual at a time when male homosexuality was illegal in the UK and often associated with "effeminacy".

> any given machine, but then again there might be other machines cleverer again, and so on.

We think that there is nothing of susbtance in Penrose's argument which was not covered by Turing's response.

More recently, Searle framed his “Chinese Room” objection to the Turing Test. He postulates a closed room into which questions written in Chinese are passed and from which answers written in Chinesen are returned. Inside the room is a non-Chinese speaker who understands English, who uses a set of rules written in English, containing pictures of Chinese symbols, to answer the questions.

3.1. **Searle and the Systems Response.** Consider the response that Searle offers those of his critics who tried the 'systems reply' to the Chinese Room argument. These critics suggested that even though Searle himself, ex hypothesi, does not understand Chinese, nonetheless the whole Chinese-language-processing system of which he imagines himself to be part might be credited with understanding. Searle has two responses. First, he finds this reply totally implausible: Why, he asks, should the addition of ancillary apparatus which in itself is clearly not a locus of understanding, alongside the uncomprehending manipulator of this apparatus, make it any more plausible that understanding is really present? This response is something of a slippery slope. Surely any materialist (and Searle claims to be one) is going to have to conclude that even in cases where true understanding is present this is the resultant effect of the operation of sub-systems not themselves possessed of understanding. (Do brain cells 'understand'? Neurons? Carbon atoms?) But in Searle's response here he seems perilously close to committing himself to the position that if the sub-systems don't understand, neither can the whole system.

Second, Searle proposes to knock out the systems reply with a slight modification of his thought experiment. OK, he says, in case you're inclined to believe that Searle-plus-ancillary-apparatus might be possessed of understanding, let's just eliminate the ancillary apparatus. "Let the individual internalize all [the] elements of the system" (all the data banks, tables of rules, etc.). So now we are to imagine Searle able to produce perfectly adequate responses to Chinese questions on Chinese stories without the benefit of his books of rules, filing cabinets full of Chinese symbols and so on – all he needs is in his head. By hypothesis, he doesn't understand a word of Chinese, but is this hypothesis still in good standing, is it still intelligible? Certainly, he could memorize some stock responses to some stock questions and that wouldn't amount to 'understanding Chinese', but then it wouldn't get close to passing the Turing test either.

Suppose Searle had started out his Chinese Room argument this way: "It's obvious that understanding can't consist merely in following the right program, even if the program is able to duplicate the observable manifestations of understanding. After all, you can imagine me speaking perfect Chinese, following the right program to 'match' my output-noises with the Chinese input-noises I hear, while not really understanding a word of the language." Would we be inclined to grant this as a starting point?

And if not, would it really help if he said, "Of course, if you have a problem with imagining that, just think of it as the internalization of a system in which I have all these filing cabinets full of Chinese characters, plus this huge rule book and... etc."? (For more on this, see the Dialogue below.)

3.2. **Searle and the Other Minds Argument.** Consider also Searle's response to the 'other minds reply' (i.e. the argument, If you're not going to credit the Chinese Room with understanding, how can you be sure that anybody apart from yourself understands anything – since all you have to go on is behavior in any case). Searle brushes off this point rather too lightly. If we take seriously his suggestion that a person might flawlessly 'simulate' an understanding of Chinese, how indeed can we know that all of our fellows are not such flawless simulators?

One of Searle's recurrent themes is the idea of the 'causal powers of the brain'. Understanding, he says, cannot merely be the result of implementing the right program; rather it can be produced only by something with the same causal powers as the human brain.

Perhaps other physical and chemical processes could produce exactly these effects; perhaps, for example, Martians also have intentionality but their brains are made of different stuff. That is an empirical question, rather like the question whether photosynthesis can be done by something with a chemistry different from that of chlorophyll.

But how could this be an empirical question for Searle?

How could we distinguish empirically between extraterrestrials endowed with genuine intentionality, and extraterrestrials who merely duplicate the outward manifestations of intentionality (pass the Turing test, etc.), while being engaged internally in nothing more than the manipulation of uninterpreted symbols?

On the other hand, any process that produces all the observable effects of photosynthesis just is photosynthesis, whatever its chemistry may be.

3.3. **Rumours of Cartesianism.** A further aspect of Searle's argument deserves attention. It would seem to be a natural response to Searle to suggest that his position – by severing 'true understanding' from its behavioral manifestation – involves a kind of Cartesianism: understanding becomes something mysterious and hidden. But Searle attempts to turn the tables on his 'strong AI' opponents by accusing them of a subtle form of Cartesianism. Consider, he says: we are not tempted to confuse such definite material processes as fires or digestion with computer simulations of these processes. Neither, he argues, should we confuse a computer simulation of understanding with the real thing. The strong AI position, however, seems to involve the claim that understanding is unique: it alone, it appears, can be detached from its concrete material implementation (in actual human brains), for its simulation is taken to be the same as the real thing. The claim that understanding or intentionality can be implemented in any medium capable of realizing the requisite program is itself, Searle suggests, a form of Cartesianism, in that it grants a special privilege to the mental.

This argument is more ingenious than persuasive. Take the case of the computer simulation of the process of digestion of a ham sandwich. The inputs to this process are presumably formal-structural descriptions of the ham sandwich and the various enzymes, etc., that will work upon it; while the outputs are formal descriptions of the sugars, amino acids, etc., that are obtained via the digestive process. In the case of the actual digestive process, the inputs are the ham sandwich and the enzymes, and the outputs the actual products of biological decomposition. Clearly, nobody will be tempted to confuse a description of a ham sandwich (however mouthwatering) with the genuine article.

Now consider the process of understanding (a story, an argument, a news item, a conversational gambit). Here the inputs are noises in the air, marks on paper, patterns on a video screen. And the characteristic observable 'outputs' of such understanding are again noises or marks of one kind or another. Furthermore, if I am asked to demonstrate my understanding, the physical medium in which I take the questions and the physical medium in which I register my answers are - within certain limits – a matter of indifference. Now the point is that in a 'computer simulation' of understanding, the characteristic inputs and outputs are (or could easily be) just the same as in the case of understanding by human beings. It's not that humans take in, and put out, real words, while computers merely take in and put out formal descriptions of words. Both entities take in and put out patterns implemented in various diverse physical media, which may be interpreted as words. Searle might object here that only human brains, or objects with equivalent 'causal powers', can 'really' put a verbal interpretation on such patterns. But that is surely to beg the question, and anyway it breaks the analogy (or disanalogy) that Searle is trying to draw: in distinguishing actual digestion from a computer simulation of digestion, it is not necessary to maintain that only human stomachs (or organs with equivalent causal powers) are really capable of interpreting slices of meat inserted between slices of buttered bread as ham sandwiches. Ham sandwiches and formal descriptions of ham sandwiches are utterly and uncontroversially different things, while marks on paper made by a human at a typewriter and marks on paper generated by a computer program may be quite indistinguishable.

Given that both people and hypothetical candidates for the Turing test, whatever their internal architectures, deal in the same range of typical physical inputs and outputs, it is as if, by way of analogy, our digestion-simulator actually took in ham sandwiches and produced sugars, amino acids and so on. But then wouldn't we speak of 'artificial digestion', rather than merely 'simulated digestion'? (Although we'd probably withdraw this designation if we discovered some simple trick, e.g. the ham sandwich just goes into a disposal unit, while the sugars etc. are merely dispensed from a pre-existing store inside the machine.) The only way out for Searle, it would appear, is to maintain that the characteristic 'output' of the process of understanding is not a set of observable responses (verbal and/or behavioral), but rather an intrinsic internal state of understanding. In that case, one might argue that there is no reason to suppose that a computer with a very different internal structure from that of the human brain would possess this sort of state, even if it produced the right sort of observable responses. This claim, however – whatever its merits – is clearly presupposed rather than established by Searle's argument.

It does not appear that the strong AI enthusiast need be unduly worried about his materialist credentials on account of Searle's charge of Cartesianism. Clearly, the former does think of mentality as something special, something whose formal properties are paramount, and whose characteristics are therefore relatively independent of their precise physical implementation. Nonetheless, of course, they can't exist at all without any physical implementation, and not any old physical medium will do.

Besides, who ever said that materialists have to hold that there's nothing special about mentality or consciousness? All they're obligated to do is show how that specialness is a function of a particular form of organization of matter. As for Searle's own position, there is clearly nothing anti-materialist or Cartesian in emphasizing the particular 'causal powers' of the human brain, as it has actually evolved and made of the stuff it is actually made of. If, however, those powers are said to include the production of an intrinsic 'internal' state, namely intentionality, the nature of which is entirely unknown to physical science then – true or false, Cartesian or non-Cartesian – it is hard to see how the position can be called materialist, at any rate as the term is typically used in modern philosophy of mind.

A Dialogue. The scene: After the Turing test. For convenience, let us say that a Chinese Searle has pulled the wool over Turing's eyes in a pseudo-English conversation. Delighted with the success of his deception, Searle cannot resist saying (through an interpreter, of course) "I really had you fooled there, Turing, didn't I?" The conversation then resumes in English, with the Chinese Searle consulting his internal look-up table for his responses.

Turing: What do you mean – had me fooled?

Searle: Well, you gave me all those complicated stories, asked me questions on them and so on, and I appeared to be understanding them.

Turing: You mean you didn't really understand?

Searle: Not a word! You see I don't actually speak any English – not in the proper sense. I've just managed to memorize a tremendously complicated algorithm for matching up English input symbols and English output symbols.

Turing: But what about right now? You seem to be conversing comprehensibly enough, although strangely, I'll admit.

Searle: My dear Turing, I haven't the foggiest idea what we're talking about. Remarkable, isn't it? This algorithm I've internalized is doing all the work, producing responses that mesh with your side of the conversation, but let me assure you: I don't understand English at all. It's all just uninterpreted symbols to me.

Turing: This is very odd. My evidence for your understanding is no weaker than my evidence for any other English speaker I've encountered. What can be missing in your case? It seems to me that your algorithm – whatever it is – does amount to understanding.

Searle: What's missing? Only the most important thing of all: real, internal, understanding. Intrinsic intentionality, my dear Turing, is what is entirely lacking on my side of the conversation, so long as it's conducted in English.

4. Centralisation

> 'The same argument (not similar but same) that cuts against the case for conceiving of the economic problem that society confronts as one in principle capable of solution through centralised planning also cuts against concieving of the mind as a hierarchical system under the command of a single unifying will.' (page 6)

The view that the mind is not a system under the command of a single unifying will is certainly something we would go along with, as would most oponents of Searle such as Dennet (Dennet 1980, Dennet 1991)or the Churchlands (Churchland and Churchland 1998).

This point is essentially irrelevant to the argument between Searle and others over the possibility of artificial intelligence. Nor is the notion of a single unifying will and or 'central' planning essential to the proposals of Cockshott and Cottrell. The representation of socialism as being under the control of a single will, the fat director theory of socialism, derives from von Mises(von Mises 1935, von Mises 1949) not Marx. Nowhere, does Marx talk about a future socialist economy being subject to a 'single will'.

What Cockshott and Cottrell argue for is planning under democratic control using modern computer technology. Whether this is implemented using central supercomputers or a distributed network of local machines, or some combination of these is an essentially pragmatic issue relating to the technology available. There are however a number of practical advantages from the centralisation of certain computation and control facilities.

The speed with which a complex decision making apparattus can function depends both upon how fast information can propagate through it, and upon how fast its individual components can respond to this information. One of the arguments against the market is that the price signals it transmits have, except in financial markets a relatively slow rate of propagation. This is because changes in price come about through changes in production and their frequency is bounded by the rate at which productive capacity can be adjusted. This implies a relatively long, and very costly, cycle time - we typically measure the business cycles as having a duration of 3 to 7 years. In contrast a cybernetic planning system could work out the intermediate and capital goods implications of a change in consumer demand in hours or days. Just how fast it would work would depend on whether the calculation used distributed or centralised computing techniques.

One component of a cybernetic control system has to be distributed. Clearly it is the Airbus factories that have the information about what parts are used to make an A340, the car plants have the information about what parts are used to make a Mondeo. This information approximates to what Boettke and the Austrian school of economics call contextual knowledge - but it is of course no longer human knowledge. Literally nobody knows what parts go into an A340. The information, too vast for a human to handle, is stored in a relational database. At an earlier stage of industrial development it would have been dealt with by a complex system of paper records. Again the knowledge would have been objective - residing in objects rather than in human brains. The very possibility of large scale, co-ordinated industrial activity rests upon the existence of such objectivised information.

The information to construct the parts explosion is generated by a computerised design process within the collaborating factories of Airbus Industrie. In a cybernetically controlled socialist economy, the parts explosion data for the A340, along with the parts explosion data for other products would have to be computationally combined to arrive at a balanced production plan.

This computation could be done either in a distributed or a centralised way. In the one case it would proceed by the exchange of messages between local computers, in the other, the parts explosion data would be transmitted to a single processing center to be handled by highly parallel super-computers.

If one uses widely distributed parallel processors the speed of computation tends to be markedly slower than when one uses tightly coupled parallel machines. If the computation requires extensive inter-communication of information - as those involved in economic equilibration do, then it becomes bounded by the transmission speed of messages from one part of the computational system to another. A tightly coupled computing system with n processors will tend to compute faster than a distributed system with n equivalent processors. This is because the communication channels between processors are shorter in the tightly coupled system, and in consequence messages travelling at the speed of light pass between processors in less time. A cybernetic system of economic control using computer technology will faster than a market one, since the electronic transmission of messages between computing centres is orders of magnitude faster than a process of price adjustments brought about by overshooting or undershooting demand; but because of the light speed limit on electronic messages there are advantages to centralising part of the computational process in the cybernetic system.

Boettke argues against the unity of the human conciousness and the concept that we have within us a single directing will. This view is in line with the research program of computational neurscience which emphasises that brain processes operate by a process of vector to vector transformation effected by a massively parallel synaptic matrix[3]. But although neural systems are highly parallel, those with the highest degree of parallelism - the brains of some families of mamals, also show high degrees of centralisation. We have a *central* nervous system, which one can contrast to those of, for example, millipedes. Millipedes have a distributed nervous system with segmental ganglia responsible for controlling the 4 legs in each segment. They also have a somewhat restricted behavioural repertoire compared to higher vertebrates.

It would, in principle for an animal like us to have a decentralised nervous system with a large number of distributed ganglia spread about the body. The cost of doing this would be slower reaction times, greater axonal mass, and greater physiological energy spent transmitting neural pulses. These evolutionary pressures have worked against the development of sophisticated behaviours in organisms whose *bauplan* involves a distributed nervous system.

Boettke argues that with the human nervous system:

> We do not build up our picture of the world from isolated facts that are placed together and then given meaning and significance. Humans do not proceed, as machine intelligence seems to by necessity do, from an isolated part to a whole. Instead, human intelligence seems to move from a grasping of the whole and then when necessary analyzing the isolated parts in light of the whole.

As a description of how visual perception works, this is onesided. It would be impossible for the visual system to grasp the whole unless it had first taken in the component details. Increasing experimental data indicates that the initial process of visual perception actually does work bottom up. The retina operates using center-surround inhibition to pick up tiny contrasting components. The primary visual cortex encodes regions of the visual field into local features, for example edges at particular gradients orientations or spatial frequencies. These then pass through successive layers were higher order statistics or properties are detected.

Boettke and Subrick's claim that machine intelligence must work in an exclusively bottom up fashion is also wrong. Computers are universal information processors, they can execute either bottom up or top down recognition algorithms. Recent progress with machine vision has come through a combination of this bottom up processing with an attentional focus to direct the gaze. By modelling the

[3]Among the foremost exponents of this view are the Churchlands whose book On the Contrary, is both an exposition of connectionism and a polemic against Searle.

primate visual system (Balasuriya 2006), and projecting the high order statistics of local features into what is called Hough Space, it becomes possible to automatically recognise the features of an object that are invariant under spatial translation and rotation.

Subjectively it may seem that we just become aware of an object as a whole rather than its low level visual components. But the fact that the lower stages of visual processing take place unconciously does not make them less real.

> Most, or nearly all of the cognitive processing has already been done when objects in concsious terms emerge, and are available for motor manipulation. Somehow, we are too late to see *how* it happened. We only notice *what* has happened.
>
> (?) page 115.

5. Motivation

> Since the fall of the communism, the former Soviet Bloc countries have had an extremely difficult time moving towards a market economy.

So Boettke and Subrick start section 4 of their article. With this we find the motivation for their paper. They are concerned with explaining this 'extremely difficult time'. The western economists who had criticised the socialist system as inefficient had anticipated that the inauguration of a market economy would lead to accelerated economic growth in the USSR. Instead it regressed from a super-power to an economic basket case. It became dominated by gangsterism. Its industries collapsed and it experienced untold millions of premature deaths, revealed in the statistics of a shocking drop in life expectancy.

A discipline less sure of itself than economics, might question its starting hypothesis when an experiment went so drastically wrong. Boettke and Subrick instead attempt to use the Searlean distinction between syntax and semantics to explain this signal failure of economic advice. They claim that the shock therapy in the USSR had changed the syntax of the economy but not the semantics: "Just because the political structure collapsed, there is no reason to assume that the social structure did. Social arrangements persisted prior to and after the fall of communism. The reformers and western advisors failed to acknowledge that the newly freed countries were not tabula rasa. They were instead countries that had residents who held beliefs about the world and the structure of society." These beliefs and attitudes that persisted from socialism are then blamed for the economic collapse[4].

This is not the place to discuss the Russia's economic trajectory during the 1990s, so we will restrict ourselves to methodological issues. Is the distinction between syntax and semantics applicable to economics in this way?

Within the discipline of linguistics the term syntax has a well defined usage. Syntax refers to systems of rules that structure the production of strings of symbols. But is there anything in economics that correponds to this?

If there is not, all Boettke and Subrick have given is a literary allusion not a theory. The school of Austrian economics from which they stem has no tradition of using the apparattus of formal syntax : productions rules, term re-write rules etc,

[4]This is reminiscent of the way the 'poisonous weeds from the past' in mens minds were an explanation for economic problems in China during the Cultural Revolution.

in its study of the economy. Ironically, to find something of this sort one has to turn to Marxist economics. Presumably because of his training in Hegelian logic, Marx's Capital (Marx 1954) opens with an analysis of commodity exchange and and the circulation of capital that uses a formal apparattus very similar to that of generative grammars.

He introduces the notion of the cirucit of capital ss $M \rightarrow C \rightarrow M'$ where M stands for a quantitity of money, M' stands for an augmented sum of money, and C for commodities purchased with the initial money. If we rewrite this as

$$M \rightarrow C$$
$$C \rightarrow M'$$

then we have in Chomsky's terminology a derivational phrase structured grammar that will produce the language $M, C, M', C', M'', C'', M''', ...$ modelling the process of growth of a capital.

What Boettke and Subrick are attempt to move towards with their syntax/semantics distinction applied to a society is something very like what Marx's distinction between base and superstructure[5]. Marx was concerned from the outset with the historical process of transition between forms of economy - modes of production. In the course of producing a theory adequate to this task he had to make an epistemological break(Balibar 1978, Bachelard 1970) from the humanist conceptions of the economy that had preceeded him(Althusser and Balibar 1970). According to Althuser such an epistemological break is the process in which a new self acting material process is discovered in an area that was once explained by ideology. The Darwinian revolution reveals an autonomous material process - evolution by natural selection - that causes the phenomena that had previously been explained in terms of divine will, or the great chain of being etc.

The autonomy here is not an autonomy of the human agents involved in this discovery - they are of course influenced the the ideologies that provide their neural 'software'. It refers to the material process that the science explores. Evolution goes on independently of human intentions, and indeed independently of the existence of the human species.

Similarly, planetary and stellar motion goes on independently of the existence of humanity and accords no special place in the universe to the world on which we live, nor is it the result of the actions of gods - imaginary projections of monarchs into the stars. Scientific astronomy replaces a notion of the skies as a reflection of human social roles with a recognition of it as an impersonal process.

Applied to the economy it replaced a conceptual framework centered on the human subject with one whose object of investigation was the mode of production and the social formation. These modes of production generated both classes of economic agents and technological changes. The resulting class struggle and the supercession of modes of production took place independently of human knowledge of it, as an autonomous causal process. Human subjects and forms of subjectivity are then constituted by the roles assigned to them in the social structure. They are something to be explained rather than an assumed starting point. The social forms

[5] It might be objected that there was a metaphorical character to this distinction in Marx. So there was. But a century and more of theoretical writings by other Marxists have given a dense social-theoretical content to what were once architectural metaphors. It remains to be seen whether the Austrian school can achieve a similar theoretical development of Boettke's syntax/semantics dichotomy.

of commodity circulation are then seen as constitutive in the construction both of the juridical subject (**?**) and the psychological subject (Althusser 1971a).

The parallel between this and Turing or Dennet's theory of mind is evident. The ideologies that constitute subjectivity are the software run by the computing engines in people's heads. The contrast between such a historical materialist approach and the constitutive subjectivism of the Austrian school is equally evident.

Once the Austrian economists became proponents of social engineering, they started to encroach, albeit in reverse gear, a traditional concerns of Marxian economics: transitions between modes of production. But they approached it with a theoretical framework inimical to the object under study. Boettke and Subrick's paper having re-asserted subjectivism was reduced to metaphors borrowed from linguistics when it attempts conjunctural analysis.

References

Althusser, L.: 1971a, Ideology and ideological state apparattuses, *Lenin and philosophy*, New Left Books.

Althusser, L.: 1971b, Lenin and philosophy, *Lenin and philosophy*, New Left Books.

Althusser, L. and Balibar, E.: 1970, *Reading Capital*, New Left Books.

Bachelard, G.: 1970, *La philosophie du non: essai d'une philosophie du nouvel esprit scientifique*, Presses Universitaires de France, Paris.

Balasuriya, L. S.: 2006, *A Computational Model of Space-Variant Vision Based on a Self-Organised Artificial Retina Tessellation*, PhD thesis, University of Glasgow Dept of Computing Science, Glasgow, Scotland.

Balibar, E.: 1978, From bachelard to althusser: The concept of 'epistemological break', *Economy and Society* **7**(3), 207–237.

Boettke, P. J. and Subrick, J. R.: 2001, From the philosophy of mind to the philosophy of the market, *CAHIERS D'ÉPISTÉMOLOGIE* **2001-03**(276), 3–18.

Chomsky, N.: 1956, Three models for the description of language, *IRE Trans. on Information Theory* **2(3)**, 113–124.

Churchland and Churchland: 1998, Could a machine think, *On the Contrary*, MIT Press.

Cockshott, P. and Cottrell, A.: 1997, Informantion and economics : a critique of hayek, *Research in Political Economy* **18**(1), 177–202.

Cottrell, A. and Cockshott, P.: 1992, *Towards a New Socialism*, Vol. Nottingham, Bertrand Russell Press.

Dennet, D.: 1980, The milk of human intentionality, *Behavioural and Brain Sciences* **3**, 429–430.

Dennet, D.: 1991, *Conciousness Explained*, Little Brown, Boston.

Marx, K.: 1954, *Capital*, Vol. 1, Progress Publishers, Moscow. Original English edition published in 1887.

Nove, A.: 1983, *The Economics of Feasible Socialism*, London:George Allen and Unwin.

Turing, A.: 1937, On computable numbers, with an application to the entscheidungsproblem, *Proceedings of the London Mathematical Society* **42**, 230–65.

Turing, A.: 1950, Computing machinery and intelligence, *Mind* **LIX**(236).

von Mises, L.: 1935, Economic calculation in the socialist commonwealth, *in* F. A. Hayek (ed.), *Collectivist Economic Planning*, London.

von Mises, L.: 1949, *Human Action*, Hodge and Company, London.