Anomalous Monism in a Digital Universe

Introduction

Psychological explanations ('I am going to the fridge *because* I desire a beer and I believe there is beer in the fridge') are a fundamental part of our out-of-the-philosophical-armchair life: they are a reliable, effective way of making sense of the world since the dawn of mankind. Much more recently, humans developed new kinds of explanations for my trip to the fridge: for example, *computational* and *neurobiological* explanations – so the question arises naturally: how the high-level and low-level explanations can be integrated? In a nutshell, this is the "Interface Problem" (IP henceforth), identified by Bermudez (2005) as the central problem in the contemporary philosophy of psychology: 'how does commonsense psychological explanation interface with the explanations (…) given by scientific psychology, cognitive science, cognitive neuroscience and the other levels in the explanatory hierarchy?' (Bermudez 2005, p. 35).

Among the many possible answers, Bermudez – drawing mainly[1] from Donald Davidson's anomalous monism (Davidson 1970) – discusses and discard the "autonomy picture" (henceforth AP) for two main reasons: first, its inability to deliver a satisfactory metaphysics of causation for mental states (Bermudez 2005, pp. 47-49); second, since AP would make IP irrelevant, AP would also make irrelevant every scientific program pursued under the assumption that solving IP is a non-trivial, important task: 'we have not been given (…) a strong enough case for abandoning an entire research program in cognitive sciences and empirical psychology' (Bermudez 2005, p. 171).

This work reviews Davidson's *original* version of the autonomy picture – which, as we shall see, is in some respect *weaker* than Bermudez' version – in the light of digital universes, a.k.a. cellular automata (henceforth CA). CA are computational tools that (unlike standard Turing machines) can easily model AP core claims, thus providing a robust and transparent context to evaluate the arguments offered by Bermudez.

The essay is structured as follows: *Section I* defines AP, IP and the relation between the two; having introduced the main philosophical concepts, we will be in a position to sketch the general structure of our arguments. *Section III* introduces the first digital universe, so that in *Section IV* we show how it can be fruitfully used to model AP; in *Section V* and *VI* we discuss our second CA and finish our review of Bermudez' original arguments. In the end, we offer some brief concluding remarks on the methodology of *this* work.

I. Anomalous Monism vs. The Interface Problem

Davidson's starting point for AP is the idea that psychological explanations are *normative*, whereas lower-level explanations are just *descriptive*. But mental explanations are supposedly *causal* explanations, so here is where metaphysics kicks in: '(…) where there is causality, there must be a law: events related as cause and effect fall under strict deterministic laws' (Davidson 1970, p. 208). Given this principle (the so-called *Principle of the Nomological*

[1] Bermudez discusses at length also Dan Dennett's position, especially as presented in Dennett (1991). In *this* work we shall focus on the version of AP developed by Davidson, even if some of Dennett's remarks will turn out to be useful in later sections.

Character of Causality), it seems we have reached an *impasse*: strict laws are available only for sub-personal levels of explanations, so either causality or monism should go[2]. Davidson's insight is that there is no real tension between these *desiderata* (Davidson 1970, pp. 210-213). which leads us to the following constitutive claims for AP:

AP_1) Mental states are physical states.
AP_2) There are no strict mental laws.
AP_3) There are strict physical laws.
AP_4) Mental states are causally efficacious.

(AP_1) is strictly intended as a *token identity*, since – as Bermudez acutely recognizes (Bermudez 2005, p. 47) – *type identity* will not allow the required gap between sub-personal and personal level. (AP_1) gives us a naturalistic ontology of the physical world while allowing for the truth of (AP_2), presented by Davidson in the following principle: 'There are no strict deterministic laws on the basis of which mental events can be predicted and explained' (Davidson 1970, p. 208). In particular, AP (as it is usually conceived) is taken to rule out *both* mental-to-mental strict laws and physical-to-mental strict laws: in other words, i) there are no exception-less laws allowing to predict your mental states knowing your physical states (the *vertical* part of the claim) and ii) there are no exception-less laws allowing to predict your next mental state knowing your previous one (the *horizontal* part of the claim). The distinction will turn out to be important later on.

It is important to realize[3] that AP, thus defined, is somewhat a *weaker* metaphysical picture than what is discussed in Bermudez (2005): for Bermudez, the autonomous mind literally *dissolves* IP:

> The picture of the *autonomous mind* understands the mind in terms of an autonomous and independent type of explanation that has no application to the non-psychological world and that interfaces only indirectly with the types of explanation applicable in the non-psychological real. (Bermudez 2005, p. 37)

It's not just that we lack exception-less laws bridging the gap between levels: on this reading of the "autonomy requirement", the upper level 'has no application' whatsoever to understand what is going on at the lower level. It is against this stronger version of AP that Bermudez puts forward his two arguments: whether they succeed or not in their original context, our conclusion will be that they are unconvincing for the *weaker* version of AP, which is the target of *this* paper – if that is true, our version of AP may still be considered a possible solution to IP. In other words: if Bermudez' arguments work for the stronger but not the weaker version of AP, we have a reason to prefer this latter version of the autonomy picture (we shall come back to this point in the concluding section).

Before further exploring our problem, it may be convenient to test our definition in a very limited, well-known domain; let's consider the following machine:

> *S* is a digital computer working in this way: given a string of symbols generated by the user on the keyboard, *S* prints "true" on the screen when at least one of the symbols is a digit.

Somewhat inspired by the classical Marr (1982), we can explain *S* behavior at different levels:

[2] Thanks to an anonymous referee for bringing to my attention this important "switch" between the two levels.
[3] Thanks to an anonymous referee for highlighting the importance of the differences between the "original" and the "modern" version of the autonomy picture.

Functional Level) *S* is calculating the following function *f* defined over string of symbols, such that:

f(x) = 1 iff *x* contains a symbol from the set { 0, 1, 2, 3, 4, 5, 6, 7, 8, 9 }
f(x) = 0 otherwise

Algorithmic Level) *S* is running the following C#-like algorithm:

```
boolean checkDigit(string s)
{
    for (int i = 0; i < s.length; i++)
    {
        int j;
        bool result = Int32.TryParse(s[i], out j);
        if (result) { return true; }
    }
    return false;
}
```

Implementation Level) *S* is checking if some cell in such-and-such region within the physical memory layer is in such-and-such physical state.

These count as *uncontroversial* explanations of *S*'s behavior, since they all meet a very strict requirement: we can consistently predict *S*'s output starting from the input (plus, possibly, other relevant information) and applying the instructions step-by-step. Given that the Implementation Level physically constrains the others, can we claim that in this case an analogous of AP (e.g., the "autonomy of *checkDigit*") holds? Clearly *no*, since the following claims hold:

$AP\#_1$) Algorithmic states are physical states.
$AP\#_2$) There are strict algorithmic laws.
$AP\#_3$) There are strict physical laws.
$AP\#_4$) Algorithmic states are causally efficacious.

($AP\#_2$) disqualifies the autonomy of the algorithmic level in two ways: i) there are strictly horizontal laws connecting the algorithm input and output; ii) there are strictly vertical laws connecting algorithmic states to physical states (in particular, the code is *perfectly isomorphic* – through the compiler – to a machine language representation, which is then physically executed by the device).
This is interesting for two reasons, i.e. because it shows that a) it is non-trivial to find a model of AP; b) the usual cognitive sciences analogy (mind : brain = software : hardware) is not, at least *prima facie*, supporting AP.

II. The Main Argument
As anticipated in the introduction, we shall address two main arguments against AP.

First, we have the *Argument From Causality*: if AP is true, the explanation of the counterfactuals relating psychological states would indeed be reduced to causal laws entirely spelled out in the lower-level language:

> It is hard to see why counterfactuals about what would have happen were someone not to believe, for example, that Edinburgh is north of Paris should be underwritten by a causal law governing the relation between types of neurophysiological events. (Bermudez 2005, p. 48)

Second, we have the *Argument From Empirical Utility*: if AP is true, IP is dissolved, so many (apparently) fruitful scientific researches trying to bring together personal and sub-personal explanations are indeed a waste of time – there is no more an interesting empirical issue to be addressed near the Interface Problem:

> Its stress on the incommensurability of the personal and subpersonal level seems to entail that little, if anything, can be said about the relation between different levels of explanations. (Bermudez 2005, p. 49)

Our main strategy will be to show that in a world in which AP holds both arguments look unconvincing. To achieve that, we shall substitute the complex world we live in with small, metaphysically neat toy universes, i.e. cellular automata; as we shall see in the ensuing section, CA are made with a very simple recipe and still they may exhibit a qualitatively rich behavior: by framing AP in universes we can perfectly monitor (and – when needed – modify to suit our needs) we may be able to observe the relevant phenomena without confounding intuitions. In particular, these are the core steps of the argument:

1) We begin by describing a two-dimensional digital universe (the *Game of Life*) where high-level phenomena do not exhibit strict horizontal laws. (*Section III*)
2) We discuss the *Argument From Causality* in the context of *Life*. (*Section IV*)
3) Given that *Life* has strict vertical laws, we move towards a more complicated universe, a one-dimensional universe where high-level phenomena do not exhibit neither horizontal, nor vertical exception-less laws. (*Section V*)
4) We point out that the same considerations made in (2) apply here as well and discuss the empirical relevance of the Interface Problem in this context. (*Section VI*)

Now that the general strategy is on the table, it is time to work out the details: we shall come back to the debate after a small introduction to CA[4].

III. Cellular Automata, Part I

Cellular automata are '*discrete*, *abstract computational systems* that have proved useful both as general models of complexity and as more specific representations of non-linear dynamics in a variety of scientific fields' (Berto, Tagliabue 2012). CA are digital universes built with cells in a *n*-dimensional lattice: at each time step in the evolution of these universes, cells "choose" their state from a discrete and finite set of possibilities according to a *local* rule, i.e. a rule that takes into account only the states of the neighboring cells at the previous time step. For example, one rule may say that a cell should be in state S_1 at t_{n+1} if and only if at t_n three of the neighboring cells where in state S_0, and should

[4] The following section is a general introduction to CA *qua* digital universes – anyone familiar with the topic may safely skip through parts *III* and *V* to read the main features of the chosen universe. For a philosophically oriented introduction to CA, see Berto, Tagliabue (2012); the interested reader may wish to consult Mainzer, Chua (2012) for a brief survey of the research field.

be in state S_0 otherwise. The underlying dynamics of these worlds is extremely simple: given the state of the lattice at t_0, the repeated application of the rule by each cell in the universe deterministically settles the state of the world in any time step in the future.

CA are a paradigmatic example of complex systems: even if cells "decide" their next state only considering their neighborhood, regular patterns can propagate throughout the universe as if they were persistent, compact objects: 'even perfect knowledge of individual decision rules does not always allow us to predict macroscopic structure. We get macro-surprises despite complete micro-knowledge' (Epstein 1999, p. 48). While encouraging the reader to explore by herself the amazing variety of CA, we shall introduce just two worlds that will be used in the rest of the work.

The first universe is, arguably, the most popular CA in history: the *Game of Life* (Berlekamp, Conway, Guy 1982). *Life* is a two-dimensional universe, cells are square and the possible states for each cell are just { 1, 0 }. *Life* popularity is due to the incredible number of interesting patterns living in this simple world[5]. As a small illustration of *Life*'s zoology, this is the so-called *glider*:

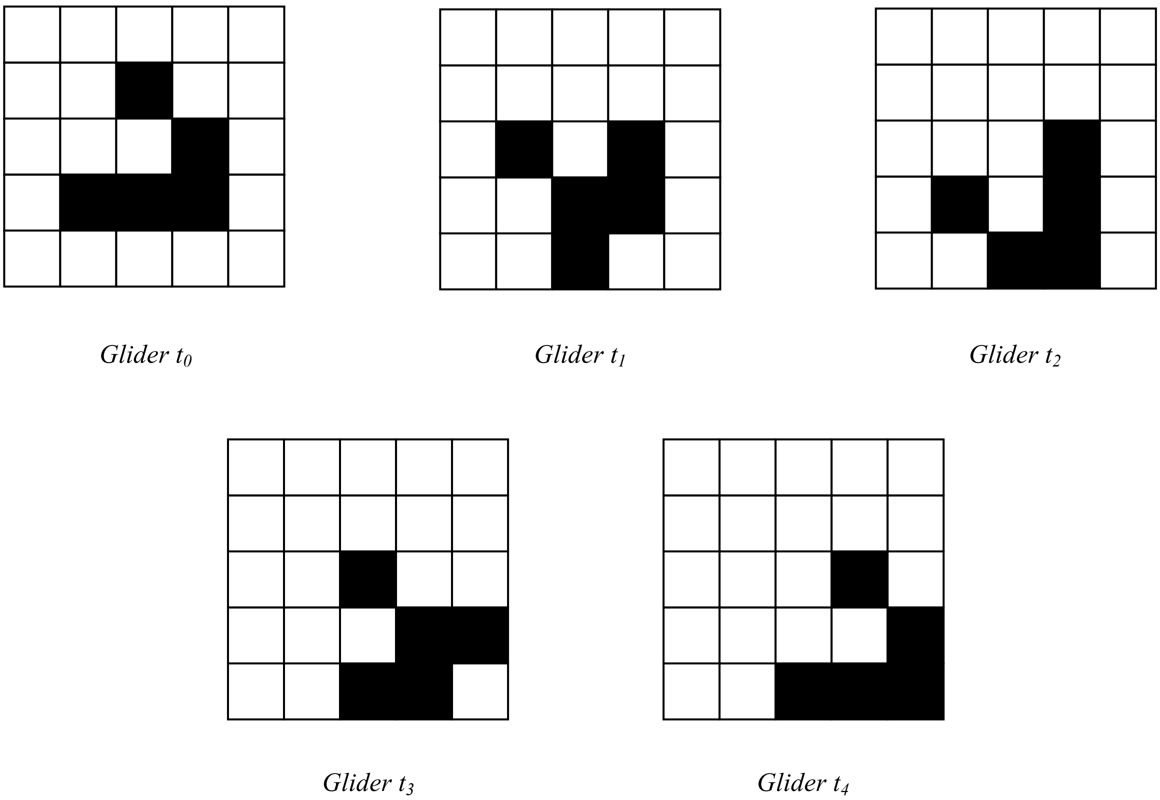

Fig. 1: a *glider* moving through *Life*.

A glider is a five cells pattern that can travel the world by repeating the above sequence of intermediate steps: while watching gliders seamlessly moving in the universe, it is important to remind ourselves that the dynamics of the universe – i.e. the only thing we actually wrote down in the code – is completely determined by the behavior of individual cells (CA are, so to speak, the paradise of reductionism – see Tagliabue (2013)).

IV. AP in a Digital Universe, Part I

[5] As a small illustration of this fact, consider that the Life Wiki (*http://www.conwaylife.com/wiki/Main_Page*) offers a downloadable package of more than 3000 of different patterns (last update, August 2013).

Now that *Life* has been introduced, the next step will be to translate AP into this digital universe. We claim that our toy universe is a model for the following AP* axioms:

> AP^*_1) Gliders are cell-by-cell aggregates.
> AP^*_2) There are no strict glider-to-glider laws.
> AP^*_3) There are strict cell-by-cell laws.
> AP^*_4) Gliders are causally efficacious.

(AP^*_2) is, of course, the crucial point. It is exactly because gliders do indeed exhibit important regularities that they are so famous among *aficionados*: if you observe the universe at t_0 as in *Fig. 1*, you can predict that at t_4 the matrix of cells will look exactly as depicted. However, the reliability of these predictions only goes that far: in a densely populated universe, it is likely that some random debris resulting from the collision of other objects will interfere with the glider's path and completely alter its future. In other words, *Life* is a nice, metaphysically "transparent" model of a world with no strict *horizontal* laws relating higher-level entities and exception-less laws for lower-level entities[6].
The *Argument From Causality* can be translated in this simple world: suppose we want to explain the matrix at t_4 given the initial conditions of the universe. At the lowest level, we have the cell-by-cell dynamics: with some patience it is easy to verify that repeatedly applying the basic law to all the cells will result in the final matrix. At the glider's level, we have some "emerging laws of motion", allowing us to predict the position of the glider without considering the underlying dynamics. This second, supposedly causal, explanation obviously supports counterfactuals: had there not been a glider at t_0, a glider would have not been in the matrix at t_4. So far, so good: AP* looks like a nice version of monism that can accommodate different levels of explanation in a neat and unified metaphysical framework. We have now two ways of interpreting Bermudez' original complaint: on the first reading (see also the comments on epiphenomenalism in Bermudez 2005, p. 49), the reason of dissatisfaction with AP* is that gliders are causally efficacious *qua* aggregates of cells and not *qua* gliders. However, this completely begs the question against AP*: if gliders are aggregates of cells, the description that picks them out is, for Davidson, completely irrelevant:

> The principle of causal interaction deals with events in extension and is therefore blind to the mental-physical dichotomy. (Davidson 1970, pp. 215).

The friend of AP* can even concede to Bermudez that, *prima facie*, cell-by-cell explanations *look* somewhat irrelevant to glider-type explanations: however, in the light of sound theoretical arguments, this is revealed to be just another case of metaphysical ignorance (compare: counterfactuals about the liquidity of water turn out to be just counterfactuals about molecular bonds of H_2O, even if, *prima facie*, facts about H_2O seem pretty irrelevant to liquidity). On the second reading, Bermudez may refer to the truth conditions of the high-level counterfactual: on this version of the argument, the problem is that the truth of the high-level statement depends crucially on our ability to reason against a given, fixed background of strict laws. On the standard treatment of counterfactuals (for example, Lewis 1973), Bermudez is right: when we go to the nearest world where the antecedent does not happen, we keep constant as many features as possible of the actual digital world, the laws of nature being among the most important. However, this complaint is irrelevant for

---

[6] As already noted, there are rules to predict the existence of gliders at *t* given the cell-by-cell matrix at *t*: therefore, *Life* supports strict *vertical* laws of the form 'if cells are in such-and-such states at time *t*, a glider exists at time *t*' and 'if a glider exists at time *t*, cells are in such-and-such states at time *t*'. We shall see in later sections an example of CA where even vertical laws fail to be strict in Davidson's sense.

the discussion of AP* (and AP), since it holds for many other theories of mind: for example, if the mind is a software running on the brain, counterfactuals about software states are indeed supported by the physics governing the hardware (obviously any computer is reliable only *insofar* the laws of nature guarantee the adequate physical behavior).

Now that we are more familiar with digital universes and AP, we shall introduce a slightly more complex universe and address even the second part of Bermudez' worries.

V. Cellular Automata, Part II

Our second digital universe (let's call it U_{17083} for reasons that shall be obvious in a moment) is a one-dimensional world whose fundamental dynamics, "Rule φ_{17083}", was discovered by James Crutchfield and Melanie Mitchell (Crutchfield, Mitchell 1995). Crutchfield and Mitchell were looking for CA rules to solve the following problem:

> *Classification Problem*) Is there a rule such that, after *n* time-steps, the world will be completely *white* (*black*) if and only if at t_0 the lattice had more *white* (*black*) cells?

As an example, consider the following universe (the vertical dimension represents consecutive time steps, being this world one-dimensional in space), generated with an initial state with more white cells than black cells:

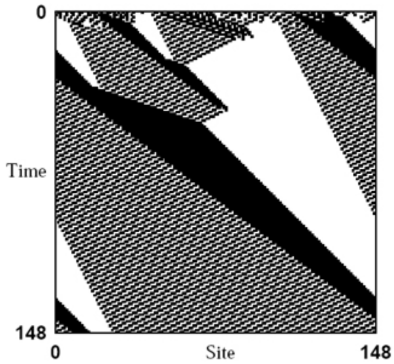

Fig. 2: space-time diagram for Rule φ_{17083} from Crutchfield, Mitchell (1995).

At time-step 250 (not shown) the grey areas disappear and all cells end up white, i.e. U_{17083} was able to solve this instance of the classification problem (obviously enough, U_{17083} can solve the problem in many instances, as detailed in Crutchfield, Mitchell 1995).

As most CA, this world is able to generate a complex global pattern: as *Fig. 2* illustrates, black, grey, white regions can be spotted in the space-time diagram. Thanks to a mix of ideas from the theory of computation and statistical mechanics (Hanson, Crutchfield 1992), the evolution of U_{17083} can indeed be mathematically "filtered" to yield a high-level representation of the dynamics (Mitchell and Crutchfield call "domains" the homogenous spatial regions in the picture, see Crutchfield, Mitchell 1995, p. 10745):

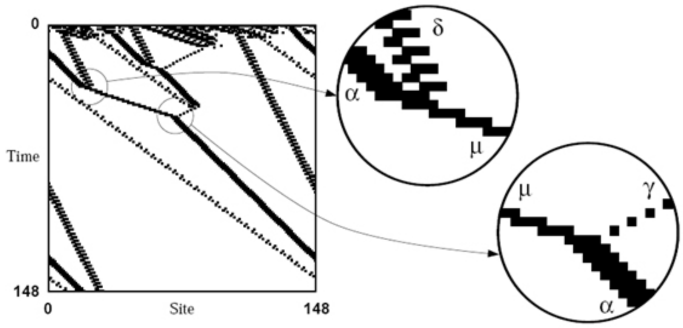

Fig. 3: filtered space-time diagram from Crutchfield, Mitchell (1995).

The new space-time diagram can now be observed as a universe whose physics is yet to be described – borrowing the words Andrew Ilachinski used for the one-dimensional CA known as Rule 110:

> Noticing that the figure consists of certain particle-like objects sprinkled on a more-or-less static background, the simplest (most natural?) thing for you to do (…) is to begin cataloging the various "particles" and their "interactions". (Ilachinski 2001, p. 662).

Crutchfield and Mitchell use exactly the same high-level language to describe their filtered automaton: in particular, when the boundary of a domain remains spatially localized over time, then they call the domain a "particle" (the long black lines in the picture can be thought as the "trajectory" of these particles):

> Embedded particles are a primary mechanism for carrying information (…) over long space-time distances. (…) Logical operations on the signals are performed when the particles interact. The collection of domains, domain walls, particles and particle interactions for a CA represents the basic information processing elements embedded in the CA's behavior—the CA's "intrinsic" computation. (Crutchfield, Mitchell 1995, p. 10744)

There are five stable particles (termed α, γ, δ, ε, μ) and one unstable particle (β): their interaction (annihilation, decay, reaction) supports the emergent logic of the system. The two circles in the image above are examples of what may happen during an interaction. In the first case, $\alpha + \delta \rightarrow \mu$, a spatial configuration representing high, low, and then ambiguous densities is mapped to a high-density signal μ; in the second, $\mu + \gamma \rightarrow \alpha$, a spatial configuration representing high, ambiguous, and then low density is mapped to an ambiguous-density signal α (Crutchfield, Mitchell 1995, p. 10745). Once this simple "algebra of particles" is obtained, it is possible to test whether this high-level dynamics can be exploited to predict the evolution of the universe without computing the future cell-by-cell: surprisingly enough, the efficiency of particles in the prediction of the final state of the universe is around 90% (Hordijk, Crutchfield, Mitchell 1996).

VI. AP in a Digital Universe, Part II

It should be obvious that the same considerations about explanations and counterfactuals from *Section IV* apply to U_{17083} as well; indeed, this second example carries arguably more weight than *Life*, considering that the distinction between levels of explanation is mathematically defined, not just the reflection of some phenomenological intuitions. It is even more remarkable that U_{17083} is a world lacking horizontal *and* vertical laws: on the horizontal side we have nice regularities (which are very far from being exception-less laws), on the vertical side, given the state of the lattice at t_x, there is no way of deducing the state of the "particles" at t_x, since it also depends on other parts of the space-time diagram. The particle-based language of U_{17083} can very well be a clear instance of 'the statistical effect of very many concrete minutiae producing, as if by a hidden hand, an approximation of the "ideal" order' (Dennett 1991, p. 43), a figure of speech introduced by Daniel Dennett, who, however, didn't provide very concrete examples (as Bermudez himself acutely points out – see Bermudez 2005, p. 142).

It is now time to address the *Argument From Empirical Utility* within our digital world: the autonomy picture supposedly threatens to make useless 'an entire research program in cognitive science and empirical psychology' (Bermudez 2005, p. 171). Again, let's try to model in U_{17083} the debate. Suppose we know nothing about U_{17083}: we look now, as for the first time, at the space-time diagram in *Fig. 2*. At first we are likely to develop a rough, but intuitive high-level explanation, in which we talk about black, grey, white regions expanding, colliding and disappearing. Now, let's suppose some smart guy among us discovers the underlying dynamics of the universe – the physics of U_{17083}. The discovery introduces a second, deeper, level of explanation for what happens in the picture, a level which, *prima facie*, is completely unrelated to what we, naïve philosophers, have observed in the first place. Truth is, in U_{17083} the Interface Problem is alive and kicking: it is indeed a profound scientific question what is the relation between the two levels – it is by trying to solve this problem that Hanson and Crutchfield came up with the idea of "filtering" the physical layer to highlight the emerging dynamics of the world. The problem Hanson and Crutchfield solved for a particular class of CA was so important that its general version was included in the seminal paper by CA pioneer Stephen Wolfram, 'Twenty Open Problems in the Theory of Cellular Automata' (Wolfram 1985); moreover, it is an important open question whether the same strategy can be generalized, for example, to "filter out" the gliders in *Life* (Miller, Page 2007, p. 233-234).

It should be clear from the above examples that the *Argument From Empirical Utility* looks rather unconvincing when modeled through a CA: while it is true that an AP-world is somewhat peculiar in its multi-level dynamics, nothing seems to imply that AP would have disastrous consequence for empirical sciences.

Conclusion

In *this* work we have argued that cellular automata can be a robust and precise model that helps us better understand the structure and the consequences of the Autonomy Picture. We have showed not only that the picture we defined is consistent, but also that it would *not* make the research agenda of cognitive sciences useless: as witnessed by the framework of computational mechanics, the very existence of autonomy does not make the interface problem irrelevant or scientifically uninteresting.

As first remarked in *Section I*, we discussed the *Argument From Causality* and the *Argument From Empirical Utility* in the context of the "weaker" version of AP, while Bermudez' original arguments were targeted against a stronger version of the autonomy thesis. Therefore, we can draw the following conditional conclusion from our discussion: *if* we have reasons to believe that Bermudez' points are actually convincing in their original setting, the arguments of *this* work strongly suggest that the best interpretation of AP is the weaker, not the stronger one.

As a final comment, we wish to point out the importance of toy universes for clearing the grounds from possible misunderstanding: while CA (even in their amazing variety) cannot *solve* by themselves philosophers' problems, modeling complex conceptual issues in *transparent*, *programmable*, *observable* universes may hopefully provide a fresh perspective on many traditional debates in the discipline (Tagliabue 2013).

Ackowledgment

The author wish to thank Francesco Berto, Edoardo Datteri, Guglielmo Feis, Luca Gasparri, and two anonymous referees for helpful comments on earlier versions of this draft.

References

Bermúdez, J. L. (2005). *Philosophy of Psychology: A Contemporary Introduction*. London: Routledge.

Berlekamp, E., Conway, J., Guy, R. (1982). *Winning Ways for Your Mathematical Plays*, Vol. 2. London: Academic Press.

Berto, F., & Tagliabue, J. (2012). Cellular Automata. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, URL = <http://plato.stanford.edu/archives/sum2012/entries/cellular-automata/>.

Crutchfield, J. P., & Mitchell M. (1995). The evolution of emergent computation. *Proceedings of the National Academy of Sciences*, 92(23): 10742.

Davidson, D. (1970). Mental Events. In Foster and Swanson (eds.), *Experience and Theory*. London: Duckworth.

Dennett, D. C. (1991). Real Patterns. *The Journal of Philosophy*, 88: 27-51.

Epstein, J.M. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5): 41–60.

Hanson, J. E., & Crutchfield, J. P. (1992). The Attractor-Basin Portrait of a Cellular Automaton. *Journal of Statistical Physics*, 66: 1415–1462.

Hordijk, W., Crutchfield, J. P., Mitchell, M. (1996). Embedded Particle Computation in Evolved Cellular Automata. In *Proceedings of the Conference on Physics and Computation*.

Ilachinski, A. (2001). *Cellular Automata*. Singapore: World Scientific Publishing.

Lewis, D. K. (1973). *Counterfactuals*. Oxford: Blackwell.

Mainzer, K., & Chua, L. (2012). *The Universe as Automaton. From Simplicity and Symmetry to Complexity.* Heidelberg: Springer.

Marr, D. (1982). *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: W.H. Freeman and Company.

Miller, J. H., & Page, S. (2007). *Complex Adaptive Systems. An Introduction to Computational Models of Social Life*. Princeton: Princeton University Press.

Tagliabue, J. (2013). A New Kind of Philosophy. In *Irreducibility and Computational Equivalence*, H. Zenil (ed.), Springer Verlag.

Wolfram, S. (1985). Twenty Open Problems in the Theory of Cellular Automata. *Physica Scripta*, T9: 170–183.