Modelling Inequality

Karim Thébault[*1], Seamus Bradley[†2], and Alexander Reutlinger [‡2]

[1]*Department of Philosophy*, University of Bristol
[2]*Munich Center for Mathematical Philosophy*, LMU Munich

Abstract

Econophysics is a new and exciting cross-disciplinary research field that applies models and modelling techniques from statistical physics to economic systems. It is not, however, without its critics: prominent figures in more mainstream economic theory have criticised some elements of the methodology of econophysics. One of the main lines of criticism concerns the nature of the modelling assumptions and idealisations involved, and a particular target are 'kinetic exchange' approaches used to model the emergence of inequality within the distribution of individual monetary income. This paper will consider such models in detail, and assess the warrant of the criticisms drawing upon the philosophical literature on modelling and idealisation. Our aim is to provide the first steps towards informed mediation of this important and interesting interdisciplinary debate, and our hope is to offer guidance with regard to both the practice of modelling inequality, and the inequality of modelling practice.

[*]email: karim.thebault@bristol.ac.uk
[†]email: Seamus.Bradley@lrz.uni-muenchen.de
[‡]email: Alexander.Reutlinger@lrz.uni-muenchen.de

1 Introduction

1.1 Econophysics and Its Discontents

The road that runs from physics to economics is an august and well travelled one.[1] In the modern context, we find that many of the models used in contemporary economics were originally created by physicists, and many practitioners of contemporary economics were originally trained as physicists. This is particularly true in financial economics.[2] The term 'econophysics' is, however, intended to demarcate something new and different from such inter-disciplinary transfers: rather it 'describes the present attempts of a number of physicists to model financial and economic systems using paradigms and tools borrowed from theoretical and statistical physics' (Stanley and Mantegna [2000], p. 355). Methodologically, then, econophysics is supposed to differ from conventional economic practice in that it uses the 'paradigms and tools' of statistical physics.[3] Along with such an approach comes an at least implicit criticism of the 'paradigms and tools' of conventional economics, including those with their origin in earlier periods in statistical physics. After all, if the methods of mainstream economics were not in some way deficient, why would it be necessary to call in the econophysicists? The notion is that, methodologically speaking, mainstream economics suffers from a number of defects that the practitioners of econophysics would like to correct. Most importantly econophysicists argue that the core principles and models of mainstream economic theory are not supported by the evidence of real economic data. Econophysics is predicated on the assumption that models and modelling techniques drawn from an experimentally focused and mathematically sophisticated science such as statistical physics will give new, and more reliable insights. Not surprisingly, not all mainstream economists agree with such a dismal view of their science. In particular, Gallegati *et al.* ([2006]) criticise econophysics on the grounds that econophysicists are: i) redoing work which has been done within mainstream economics; ii) ignoring rigorous and robust statistical methodology; iii) assuming universal empirical regularities where there are none; and iv) using modelling techniques that are *in certain senses* inherently problematic or illegitimate (see also McCauley [2006]). The final line of criticism is the main focus of this paper, and we will consider in particular the 'kinetic exchange' models for income distributions within populations that are the subject of most ire within Gallegati *et al.* ([2006]).

Our intended role is as a mediator rather than a champion for either side in the debate. Our hope is that by elucidating the sense in which the models in question can and cannot be taken to be problematic one can understand more about the nature of the modelling practices and idealisations used in the econophysics of social systems.[4] Our analysis begins, in Section 1.2, with a informal description of the econophysics models for income that we are dealing with. Given that, we will introduce the principal 'burglar economics' line of criticism that has been made of these models. Following this, in Section 2.1, we consider the wider context of models for income in mainstream economic theory. In Section 2.2 we introduce the first econophysics income model (the 'DY' model) in detail. We next move to isolate and evaluate the

---

[1]See Mirowski ([1992]) for an historical overview of the relationship between the two disciplines

[2]For a popular overview of the 'Physics of Wall Street' see Weatherall ([2013])

[3]As a matter of sociological fact, econophysicists typically publish in physics journals, rather than in economics or finance journals (Gingras and Schinckus [2012]). For more on the history of econophysics, see Jovanovic and Schinckus ([2013a]). For more on econophysics' status as multi-, inter- or trans-disciplinary, see Rosser, Jovanovic and Schinckus ([2010, 2013b]).

[4]For work focusing on the methodological foundations of financial econophysics see Kuhlmann ([2006, 2014]), Rickles ([2007, 2008]), Casini ([2014])

idealisations involved in the DY model, by comparison with the (remarkably similar) Boltzmannian gas model. In Section 4 we present a second, more sophisticated, econophysics income model (the 'CCM' model) based upon a specific 'de-idealisation' of the the DY model. Finally, in Section 5, we evaluate the connection between the foregoing analysis and contemporary accounts of models and idealisations in physics and social science. Ultimately our conclusion shall be rather an open one: There remain question marks regarding at least some of the idealisations involved in the econophysics models. However, it remains to be seen whether this indicates a shortcoming of the models or in adequacy of contemporary philosophical work for dealing with the modelling practice of econophysics.

1.2   Against Burglar Economics

Consider the following two economic interactions:

> *Scenario 1*: Two idealised economic agents bump into each other while carrying with them all their income[5] in cash money. All the money of both agents falls onto the floor. They then each pick up a random proportion of the money on the floor and go on their way.

> *Scenario 2*: Two idealised economic agents each put a (possibly different) fraction of their money in their back pocket. They then bump into each other and drop all the money not in their back pockets onto the floor. They then each pick up a random proportion of the total money on the floor and go on their way.

What can these types of randomised 'zero intelligence' economic interactions tell us about real economic exchange processes? Surely nothing! Except for in a few extreme scenarios, real economic interactions are almost entirely different, involving, as they do, things like expectations, credit and risk aversion. The scenarios described above bear very little resemblance to how economists usually think about microeconomic interactions.[6] Rather, they bear a very strong resemblance to how physicists think about microphysical interactions. Scenario 1, in particular, is very similar to the type of randomised binary interaction that have been used to build models of statistical systems since the mid-nineteenth century. Thus, it is prima facie very surprising that models constructed along the lines of our two scenarios have been used to recover important qualitative features of the distribution of individual monetary incomes found within real economies. In particular, when aggregated over large populations, kinetic exchange income models along the lines of the first scenario produce an exponential distribution of incomes that corresponds to the bulk shape found within real data. And, when aggregated over large populations, kinetic exchange income models with randomised savings along the lines of the second scenario, produce an exponential distribution with a power law tail in correspondence to both the bulk and tail shape of real data. These characteristic distribution shapes are found in a wide variety of real income data, including very different time periods and economies.

---

[5]Here and below we choose to focus exclusively upon income rather than principally because of the lack of easily available data source for measuring wealth. The wealth data that is available tends to be qualitatively very similar to income data, only with a lower exponent in the tail distribution, indicating higher inequality (Sinha and Chakrabarti [2012]).

[6] Zero intelligence agents are also used in economics. The goal there is to differentiate effects due to the market mechanism, and effects due to trader strategy: if a phenomenon shows up even when the agents enter into transactions at random, then the phenomenon is due to market effects only Ladley [2013]. In both economics and in econophysics, the behaviour of zero intelligence agents is not taken to represent that of real economic actors.
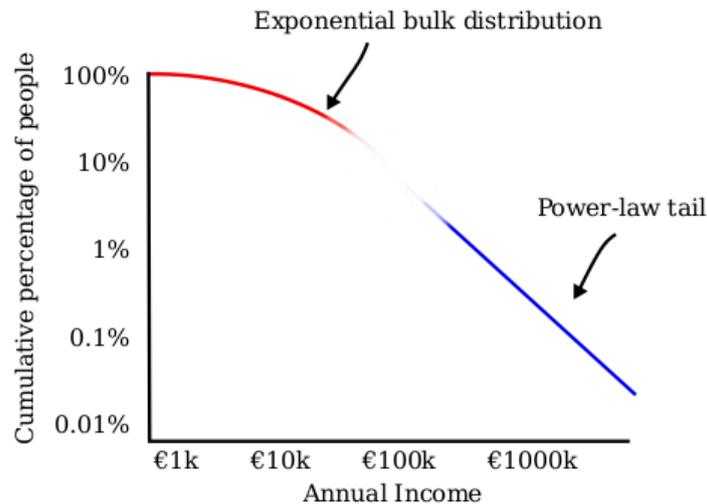
Figure 1: Illustration of 'stylised fact' of an exponential distribution with a power-law tail that is found within a wide variety of real income data. Plot is log-log.

One might, in fact, consider an exponential distribution with a power law tail (see Figure 1) as the key 'stylised fact' of income distributions, akin to those that are of central concern within the econophysics of financial markets. Stylised facts are sets of 'properties, common across many instruments, markets and time periods' (Cont [2001], p.223) and recovery the stylised fact of power-law distributions is taken to be one of (but not the only) chief achievements of econophysics. Despite this impressive correspondence with real data, kinetic exchange models of income have come under attack. Such attacks are not just from economists hostile to econophysics per se, rather some see kinetic exchange models as part of a 'worrying trend' in econophysics, that should be corrected. Such models, we are told, are 'not like [an] economic exchange process, but more like a burglar process. People randomly meet and one just beats up the other and takes their money' (Hogan [2005]). Moreover, the models involve an exchange process that:

> [Would] hardly be imaginable as both agents would rather prefer not to participate in this game of a burglar economy - at least if they possess a minimum degree of risk aversion. The model, thus, is not in harmony with the principle of voluntary participation of agents in the hypothesized process which economists would consider to be an important requirement for a valid theory of exchange activities. Lux ([2005]).

In more general and precise terms, the 'Worrying Trends' line of criticism runs as follows:

> There is a general (and understandable) tendency on the part of econophysicists to develop theoretical models, which are based on the principles of statistical physics. We obviously recognise that all theories are approximations to reality. In certain specific economic applications, models constructed on this basis may very well be valid. But in general this is most definitely not the case, and models constructed on these principles ignore absolutely fundamental features of economic reality. Any congruence obtained with the data by such models is therefore spurious. The main problems with translating statistical physics models into economics are

1. these are essentially exchange-only models of economic and financial processes, which take no account of production,

2. they often lead to a confusion of basic concepts, in particular the concepts of transactions and of income.

(Gallegati *et al.* [2006], p.4)

It is not the accuracy of the models in recovering real data that is in question: rather, Gallegati *et al.* ([2006]) are primarily engaged in a methodological attack on the idealisations involved in these models. They argue that kinetic exchange type models for income distributions are inherently problematic or illegitimate on the grounds that their treatment of production, income and transactions is in conflict with 'economic reality'. In what follows we will consider two simple but important kinetic exchange models in detail, and assess what warrant, if any, such criticisms should be given.

2 Modelling Inequality

2.1 Mainstream Economic Models for Income Distribution

The problem of modelling income distributions within societies is certainly not a peripheral issue in the history of economics, and was much discussed long before than advent of econophysics. Consider, for example, the following quote from David Ricardo,[7] one of the founders of classical economics:

> The produce of earth—all that is derived from its surface by the united application of labour, machinery and capital, is divided among three classes of the community, namely, the proprietor of the land, the owner of the stock or capital necessary for its cultivation, and the labourers by whose industry it is cultivated. But in different stages of society, the proportions of the whole produce of the earth which will be allotted to each of these classes, under the names of rent, profit and wages, will be essentially different[...]To determine the laws which regulate this distribution is the principal problem in Political Economy.

Of particular importance, from the perspective of contemporary macro-economics, are models that seek to explain the structure and evolution of income distributions in more economically developed countries in terms of the idea of 'marginal productivity' and the 'race between technology and education'.[8] The essential idea, dating back to Tinbergen ([1975]), is to consider a situation where individuals within a population have different levels of education leading to different levels of skill. As technology progresses the demand for higher and higher skill levels increases. This is because workers who perform higher skilled jobs are understood to make a greater personal contribution to the output of the company for which they work; they have higher marginal productivity. Since they are in greater demand such workers can secure higher wages than those with less education. Given that there is educational progress with a society, over a period of time the education level, and thus skill level, of the workers at the bottom will increase. If there is little technological progress for a long period of time and everyone eventually reaches a level of education sufficient to have the skills to perform any

---

[7]From the preface to his 1817 'On the Principles of Political Economy and Taxation', as quoted by Atkinson and Bourguignon ([2000])

[8]For more details on the practice of modelling income inequality within contemporary macro-economics (and micro-economics) see Atkinson and Bourguignon ([2000]) and Bertola *et al.* ([2014])

job, then the society will tend towards equality of wages. Conversely, if the pace of technological change far outstrips the pace of educational development, then the wages of a highly educated elite will become extremely high, and overall income inequality will increase. Such models make use of a range of idealising assumptions, not least the idea of competitive equilibrium markets that has been subjected to criticism from various perspectives (including from econophysics). They do, however, provide valuable insights into possible explanations for aspects of the structure and evolution of income distributions, when understood purely in terms of wages. As such it seems very difficult to consider them as entirely without value, even if they neglect important structural, institutional and cultural factors.

As emphasised in the Ricardo quote above, when we are considering income of individuals within an economy we must of course also consider income from 'capital' in addition to waged labour. The definition of the concept of 'capital' is a subtle issue. In an influential account, Piketty ([2014]) gives the definition of capital as 'the sum total of nonhuman assets that can be owned and exchanged on some market' (p.46).[9] According to Piketty ([2014]), from a historical perspective return on capital has typically occupied a significant share of total income. And furthermore, on Piketty's view, capital is of particular importance to income inequality since, in times of low economic and demographic growth, is the principal factor driving increasing income inequality. The basic idea is as follows. Consider an economy where the rate of return on capital is significantly larger than the growth rate for a long period of time (for example nineteenth century Europe). In such circumstances, if people with inherited wealth save a proportion of their income from capital, then their capital will grow faster than the economy as a whole. Overall this will mean that income from capital will grow faster than waged income and thus, since capital is very unequally spread within the population, the distribution of income will tend to diverge. Such focus upon the historical context and macroeconomic mechanisms of income inequality seems ideally suited to addressing particular culturally and historically contingent questions. These are not, however, the only questions one would like to ask about the distributions of incomes.

In particular, irrespective of the nuances behind the income distributions that obtain in particular countries in particular epochs, it is still legitimate to try and find general features of economic interactions that explain (relatively stable) qualitative features such as the 'stylised fact' discussed above. The important point, that we will return to in Section 5, is that the idealisations one makes in order to to construct a socioeconomic model (as well as the manner in which these idealisations are justified) should be evaluated relative to the type of purposes to which one wants to put the model. If one wants to understand the dynamics of the income distribution in a particular country over a period of several decades with a view to motivating policy that reduces inequality, then idealising away population growth and production is clearly entirely inappropriate. Moreover, in such a context, recovery of qualitative 'stylised facts' is of marginal importance:

> When we say that a distribution of wealth is a [power-law] distribution, we have not really said anything at all. It may be a distribution in which the upper decile receives only slightly more than 20 percent of total income (as in Scandinavia in 1970-1980) or one in which the upper decile receives 50 percent (as in the United States in 2000-2010)...In each case we are dealing with a [power-law] distribution, but the coefficients are quite different...When we study inequality in historical perspective, the important thing to explain is not the stability of the

---

[9]This definition excludes 'human capital' which includes an individuals skills and training. The adequacy and consistency in application Piketty's conception of capital has been contested. See for example Homburg ([2015]).

distribution but the significant changes that occur from time to time. (Piketty [2014], p.368)

This point not-with-standing, the recovery of the coefficients of power-law distribution also requires the recover of the distribution, and so focus upon the income distributions relevant to specific historical circumstances rather than 'stylised facts', does not reduce the explanatory burden vis-à-vis the power-law. Moreover, when what we are interested in is the question of why the tail of income distributions follows a power-law stably across time and space, then it does seem appropriate to use a highly idealised framework where we do not attempt to model the factors that determine the relevant coefficients.[10] The question of whether or not such models can be usefully applied in the context of designing inequality-reducing policy will be considered in Section 5

To conclude our discussion of income models within mainstream economics we will briefly consider a family of stochastic models that recover the power-law behaviour in stable income distributions. Most significant is the model of Champernowne ([1953]).[11] This model is built within a dynamic version of the general equilibrium framework and has the feature that the entire distribution tends towards a power-law distribution under certain fairly general conditions. Champernowne-type models bear a strong resemblance to the kinetic exchange models that are the focus of our analysis. In particular, the stochasticity that they incorporate is primitive, and not related to individual optimising behaviour. As we saw in the previous section, such features have been harshly criticised in the econophysics case on the grounds that they conflict with basic premisses of economic theory. Similarly, we find comments on the Champernowne approach such as:

> From the economist's point of view, perhaps the most unsatisfactory feature of the stochastic models...is that they shed no light on the economics of the distribution process....it is difficult to see how the factor of individual choice can be disregarded in analysing personal income distribution. (Mincer [1958], p. 283)

These remarks are very similar to the 'burglar economics' line of criticism that we saw levelled at the kinetic exchange models, only presented in a more moderate manner. For this reason we will not here conduct any detailed analysis of the stochastic models of income in economics and their relation to kinetic exchange models in econophysics. The reader is referred to Chakrabarti *et al.* ([2013])[§7.2] for some efforts in this regard.

2.2   Econophysics Models for Income Distribution

There is a suggestive analogy between the emergence of stable income distributions in economic systems, and the emergence of stable kinetic energy distributions in statistical mechanical systems. Yet it is not, at first hand, at all obvious that such an analogy should be extendable to a deep structural similarity between the models that can accurately describe such systems. There is some pedigree to such an expectation, however. For instance, we can find the great mathematician Benoit Mandelbrot (Mandelbrot [1960]) opining:

---

[10] This is a general strategy: if a phenomenon occurs in multiple contexts, your model ought not appeal to details specific to a particular context. Thus, more idealised, more abstract models are useful when the phenomenon is multiply realised.

[11] See also Wold and Whittle ([1957]), Brown ([1976]) and Chakrabarti *et al.* ([2013], §7.2). Economics model involving power laws in other contexts are considered in Kesten ([1973]), Gabaix ([1999]), Gabaix *et al.* ([2005]).

There is a great temptation to consider the exchanges of money which occur in economic interaction as analogous to the exchanges of energy which occur in physical shocks between gas molecules. In the loosest possible terms, both kinds of interactions *should* lead to *similar* states of equilibrium. That is, one *should* be able to explain the law of income distribution by a model similar to that used in statistical thermodynamics.

What happens if we take Mandelbrot's intuition at face value, and simply press ahead in an attempt to explain 'the law of income distribution' using 'a model similar to that used in statistical thermodynamics'? A surprising, and relatively recent, achievement of econophysics is the construction of viable 'gas-like' models for income distribution. Here we will introduce the most basic such model, the 'DY' model of Drăgulescu and Yakovenko [2000], and assess the relationship with its statistical mechanical analogue. In Section 4 we will consider more sophisticated kinetic exchange models in the context of the 'stylised facts' found within real income data.

The starting point for the DY model is a population of 'zero-intelligence' agents. These agents have a single property: their money. At any given time t an agent i has associated with them a single property, their monetary income $m_i(t)$ (which is always non-negative, so debt is not allowed). In the DY model one first assumes a large population (i.e. N agents, with $N >> 1$), and then randomly selects two individuals at some time t. For our two chosen agents, the initial pre-interaction state can be characterised completely in terms of two numbers: $m_i(t)$ which is the income of agent i at time t; and $m_j(t)$ which is the income of agent j at time t. Following our guiding Mandelbrotian intuition, we model all interactions in the population in terms of binary exchanges of money in the same way as in the kinetic theory of gases one can treat the interaction between molecules in a gas in terms of binary exchanges of kinetic energy.[12] However, unlike in the kinetic theory case we do not have dynamical principles dictating how the initial conditions of the agents determine the nature of the collision. Rather in the agent-agent 'collision' all the money of the two agents is pooled, and then a random fraction is given to one, and the rest to the other. This simple exchange mechanism thus leads to a post-interaction state characterised by:

$$m_i(t + 1) = m_i(t) + \Delta m \tag{2.1}$$

$$m_j(t + 1) = m_j(t) - \Delta m \tag{2.2}$$

where

$$\Delta m = \epsilon_{ij} m_j(t) + (1 - \epsilon_{ij}) m_i(t) \tag{2.3}$$

with ϵ_{ij} a random variable uniformly distributed between 0 and 1, and labelled by the index of the two agents in the interaction (i.e. agents i and j). Although the model only includes binary interactions, at each time-step we resample the two agents that are interacting from our population. This, together with the randomisation of the fractional monetary exchange, means that at late times, $t \to \infty$, the distribution of incomes within the population should be 'well-mixed'. A crucial feature to note is that in this model both the total number of agents, N, and the total amount of money, $M = \sum_i m_i(0)$, are held fixed.[13] This is in exact correspondence to the conservation of number of molecules and total energy in the kinetic theory of gases. Following Banerjee and Yakovenko ([2010]), we can characterise a

---

[12]The comparison that is particularly pertinent here is with Boltzmann's original 1872 model Boltzmann ([1970]).

[13] That is, $\sum_i m_i(0) = \sum_i m_i(t)$ for all t.
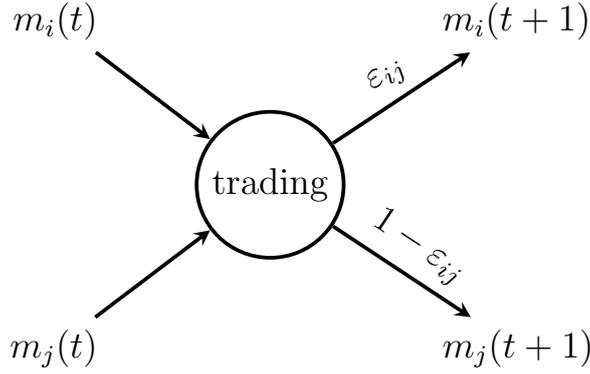
Figure 2: The DY exchange dynamics (after Chakrabarti et al. [2013], Figure 4.1, p.56)

probability distribution for money using bins of width $m_\star$ labelled by the index k. We let N_k be the number of agents with income between m_k and $m_k + m_\star$, so the probability for an agent to be in the kth bin, is just $P(m_k) = \frac{N_k}{N}$. Holding the resolution $m_\star$ fixed, the income distribution is thus characterised entirely by a set of occupation numbers, N_k for the bins. The simple combinatorial expression for the number of distinct placements of the agents that preserves the same set of occupation numbers is:

$$\Omega = \frac{N!}{N_1!N_2!N_3!...} \tag{2.4}$$

We can then *define* entropy as the natural logarithm of this quantity, $S = \ln\Omega$, and then look for a distribution that maximises entropy. The origins and status of such a principle of maximum entropy is a subtle and complex issue that would require a lengthly treatment to be dealt with adequately. Here we will simply note there is a general correspondence between the probability distribution which maximizes the entropy and the frequency distribution which can be realized in the greatest number of ways Jaynes ([2003]). For further details, on the relationship between entropy maximisation and stable probability distributions the interested reader is referred to the discussions of Frank ([2009]) and Lyon [2011, 2013].

The definition of entropy as $S = \ln\Omega$ reflects this correspondence. A large N approximation together with a little analytical work Banerjee and Yakovenko ([2010])[p.4] then leads us to the distribution:

$$P(m_k) = \frac{N_k}{N} = e^{-\frac{(m_k - \mu)}{T}} \tag{2.5}$$

with $T = \frac{M}{N}$ and $\mu = -T \ln \frac{T}{m_\star}$. This is of course identical to the Maxwell-Botzmann speed distribution for molecules in an ideal gas with, in that case, T the temperature and μ the chemical potential. The same distribution is produced at late times via Monte Carlo simulations (Chakrabarti *et al.* [2013]) given a suitably large N, implying that the model is in fact entropy maximising on a suitable timescale. So this simple model captures the stylised fact that the bulk of the income distribution is an exponential distribution.[14]

---

[14]Gallegati *et al.* ([2006]) note that there is some controversy over whether the bulk of the income distribution is exponential or lognormal. Wiens ([1999]) points out that the gamma distribution (of which exponential distributions are a special case) can also look a lot like a lognormal in some circumstances, and so this is a difficult thing to determine. Chakrabarti *et al.* ([2013]) seem to talk as if the choice of one or the other of these distributions is almost a matter of discipline-specific preference (p.2): they claim that economists prefer the lognormal, while statisticians and physicists prefer gamma distributions.

3   Idealisations in Kinetic Exchange Models

We have just seen that by taking the kinetic theory analogy seriously, and constructing an income exchange model along very similar lines to kinetic exchange models, we can derive a distribution for incomes of the same exponential form as that used in the paradigmatic model of gases. The comparison between kinetic exchange models of income and gases in fact proves to be particularly enlightening. In both cases what we are dealing with are heavily idealised models. Intuitively speaking, an idealised assumption is a 'literally false' assumption about the target system. It is literally false because, for instance, it leaves out a feature that the target system actually has, or it ascribes a feature to the target system that it actually lacks. Various accounts of idealisation exist within the philosophical literature (Ladyman [2013]). For the time being we will, so far as is possible, remain neutral between these accounts and focus upon (prima facie) literally false assumptions that require special justification. In Section 5 we will return to question of the interpretations of idealization and consider a number of contemporary accounts.

	Gas	Income
Feature 1	Multi-molecule interactions	Multi-agent interactions
Idealisation 1	Binary interactions	Binary interactions
Feature 2	Molecules have chemical properties	Agents have memory and expectations
Idealisation 2	Molecules are hard spheres	Agents have zero intelligence
Feature 3	Conservation of total energy	Production of money
Idealisation 3	No idealisation necessary	Conservation of total money
Feature 4	Conservation of total molecule number	Growth of population
Idealisation 4	No idealisation necessary	Conservation of total agent number
Feature 5	Not all kinetic energy is exchangeable via linear momentum transfer	Not all money is exchangeable via random redistribution
Idealisation 5	All kinetic energy is translational	All money is exchangeable via random redistribution
Feature 6	Molecules have intrinsic differences besides their speed	Agents have intrinsic differences besides their income
Idealisation 6	Molecules are identical	Agents are identical

Table 1: Comparison of Idealisations in Gas and Income Models

If we consider in particular the 1872 Boltzmann kinetic model (Boltzmann [1970], Emch and Liu [2002], Uffink [2014]) we can make a specific comparison between some of the key idealisations involved in income and gas kinetic exchange models. This list given in Table 1 is by no means meant to be exhaustive. Rather its function is to illustrate that in both cases idealisations relating to specific features of the real systems are made, and in both cases these idealisations vary in their type and degree of justification. Moreover, although not complete, our particular selection of idealisations has been engineered to enable a comparison between

the allegedly problematic idealisations of kinetic exchange inequality model, and those in the kinetic theory of gases which have widespread support.

3.1 Binary Interactions

Arguably the most important idealisation in both cases is that all interactions are binary. In the gas case in particular we find statements such as: 'The legitimate neglect of all but binary encounters in a gas is one of the important simplifications that have enabled the theory of gases to attain its present high development' (Chapman and Cowling [1991], p.3). The idealisation of there being only two-particle collisions amounts to saying that interactions involving more than two particles can be neglected. Where does the legitimacy of this simplifying assumption derive from? The answer to this question is different for gases and economic exchange.

In the case of a gas we have a relatively straightforward answer. The justification is in terms of a particular density of the gas: low enough that collisions involving three or more molecules are very rare, but not so low that two molecule collisions don't happen. On the one hand, encounters with other molecules should occupy a very small, but not negligible, part of the lifetime of a molecule, meaning that binary collisions are overwhelming more likely than three or more molecule collisions. On the other hand, the dynamics of the model also requires that binary collisions within the population of molecules that make up the gas should not be so rare that the average distance travelled before a collisions (the 'mean free path') is comparable to the dimensions of the container. It is a contingent fact that real gases within the physical circumstances relevant for experiments on earth actually are in such density regimes. But given this contingent condition holds, it is approximately true that only two-particle collisions occur in gases. An additional justification of the idealisation is instrumentalist in spirit and is based on the (empirical) success of the model: without the binary interaction limit holding in at least some instantiations of the target gas system the original Boltzmannian model, as well as its more sophisticated modern cousins, would neither be particularly practically useful nor empirically adequate.

With regard to the binary interactions idealisation, income exchange models are significantly different from the gas model. The difference between the two models concerns the justification of the idealised assumptions in each model: unlike in the case of the gas model, there is no straightforward and precisely quantifiable justification of the all important binary interaction idealisation for income exchange models. If this idealisation is to be justified one would expect that the manner of this justification must be rather different. In an economic system there is no sensible analogue for the 'density' where binary agent-agent interactions dominate, and no real basis to expect that the kind of complex economic systems that the model is intended (in some sense) to represent will actually be dominated by the binary interactions.

Here we should make two important comments regarding the dialectic between the economists and the econophysicists. First, binary interactions are not among the supposedly problematic idealisations mentioned in Gallegati *et al.* ([2006]). Their critique primarily derives from a comparison between the kinetic exchange income model, and the concepts used in mainstream economics models. We think this illustrates the virtue of our approach: when studying the idealisations within econophysics models it is highly instructive to consider the idealisations made in related physics models. This first point notwithstanding, we think the case of binary interactions idealisations well illustrates the general worry that Gallegati *et al.* ([2006]) clearly have in mind: when importing models from physical science into social science, one must be aware that the justifications for the relevant idealisations will not always

(or even usually) come with them. Arguments must be given in order for the models to be 'sanctioned' in their new context (Winsberg [1999]).

So, to what extent can a binary interactions idealisation be justified for agents rather than particles? One option would be as follows: agents are unlike particles in that 'three-body' interactions can be decomposed into sequences of 'two body' interactions. In this way monetary exchanges involving three individuals might be suitably approximated by sequences of exchanges involving two agents. On the face of it, this argument seems rather plausible, and we should note that many mainstream economic models also assume all interactions are binary. However, from a game theoretic perspective, it does seem that three player interactions are importantly different from games involving only two players (von Neumann and Morgenstern [1944], e.g. chapter 5). Thus, at this stage, it is not clear that the binary interaction assumption is justified for the case of kinetic exchange income models. At the least, we lack any strong arguments that the non-binary nature of real economic interactions is 'explanatorily irrelevant' when we consider the distributions of incomes in a economy. We will return to the question of whether or not the problems in justifying the binary interaction idealisation should be taken to undermine the value of kinetic exchange models in Section Section 5.

3.2 Conservation Principles

Let's consider the conservation principles next. One of the main points of criticism made of the kinetic exchange models relates to the conservation of total money and total number of agents. This can be compared with the conservation of total energy and total number of molecules. In a gas the conservation of energy is not an idealisation: it is a fundamental physical principle. Similarly, the conservation of molecule number in a real gas can be justified based upon physical principles (e.g. conservation of mass or some form of correspondence principle). For an economic system, on the other hand, money is usually understood to be produced over time, and migration, birth and death certainly do occur. Thus, the conservation principles which could be taken as non-idealised in the gas model case, must reasonably be considered idealisations for the income exchange model. The conservation of money idealisation is of particular relevance for the dialectic between the economists and econophysicists that featured in our introduction: it is one of the 'absolutely fundamental features of economic reality' that is 'ignored' in the exchange models. The money supply is, furthermore, closely connected to economic production, and so the conservation of money is doubly unrealistic.

A seemingly straightforward justification for both the conservation idealisations in the econophysics case runs as follows: Chakrabarti *et al.* ([2013], p.56) argue that since the time scale for the agent-agent collisions (the 'timestep timescale') is much shorter than that in which economic growth or migration take place, neglecting such features is a reasonable approximation. This is despite the fact that the timescale of the distribution described by the full model run could reasonably be considered that of long time, population level effects; for instance economic growth and demographic change. This justificatory strategy invites further critical questions: (i) whether money conservation is a legitimate approximation on the timestep timescale; and if so (ii) whether such a short time scale conservation idealisation is undermined by non-conservation on the timescale of the model run.

Critical question (i). With regard to the timestep timescale idealisation and money conservation, there are some grounds to be skeptical. McCauley ([2006]) argues that money is created and destroyed rapidly, 'with the tap of a computer key', via credit, and that, on these

grounds, 'conservation of money is a silly assumption'. But are such considerations enough to invalidate the collision models built upon conservation of money? If one considers real figures for annual money growth then the overall creation and destruction of money is only, by percentage, of the order of double digit changes per year (IMF [2014]). This would average out to very small changes at the timescale of the economic interactions of real agents (i.e. minutes). Even if, strictly speaking, money is created and destroyed on the timestep timescale it does not seem an unreasonable approximation to assume money conservation on such timescales. One could understand this as a contingent fact about actual economies in the same way as we understood the contingent fact that binary interactions dominate in the gas pressure regimes of interest. However, one might instead think that it is almost constitutive of a functioning economy that annual growth is moderate: a currency experiencing hyperinflation, or hyperdeflation would no longer serve as a useful means of exchange. That said, clearly there are some real financial events that happen over relatively short timescales, and yet could dramatically affect the money supply, for example a surprise injection of reserves by a central bank. The question is then, whether the money supply is 'sticky' enough that the effects of such 'money growth shocks' still spread relativity slowly, and thus are only important over timescales large relative to the timestep timescale. These are difficult issues and it would radically overstep the remit of this paper to attempt to settle them here.[15] We will only point out that the question of whether or not the assumption of money conservation on timestep timescales is reasonable is a subtle one.

Critical question (ii). This brings us to the second part of our question, whether a legitimate timestep scale conservation idealisation is undermined by non-conservation on the timescale of the full model run. The crucial point is that by the 'timescale of the full model run' we should mean the timescale of the relaxation to the equilibrium distribution within the model, rather than the timescale of the stable income distribution data. If the timescale of relaxation is short enough that money is approximately conserved, then that idealisation within the collision model should be understood to be well justified. On the other hand, if relaxation takes enough time for the growth in money to be appreciable, then the idealisation seems dubious. This essentially amounts to the question of whether we should understand the growth of money supply as something which happens much slower than the relaxation to equilibrium in the distribution of income. The point is that, unless the growth of money supply has this property, there are good reasons to doubt the applicability of any argument towards an equilibrium distribution, just like in a non-adiabatic transformation of a gas. These are subtle and difficult questions, and again we will not attempt any kind resolution here. Rather we will simply note that the question of timescales and legitimacy of money conservation idealisations appears to be much more subtle than either of the sides in the debate appreciate. If these issues are crucial in passing judgement over these models, one way or another, then they must be considered in more detail.

One interesting line of reasoning that might allow the defenders of kinetic models of income to blunt criticisms based upon the non-conservation of money runs as follows. If it is distribution of relative incomes that is the 'stylised fact' we are trying to explain, then what really matters is relative rather than absolute monetary amounts. In that case, we could take the 'money' that is being exchanged in kinetic exchange models to really be units of relative income or purchasing power. And in that case, overall conservation of 'money' is not just a reasonable approximation, it is actually genuinely the case since the total units of relative income within the population will be conserved over time. The fact that the money supply within real economies grows would simply not be relevant: if it is relative income inequality

---

[15]See Christiano ([1994]), Walsh ([2010]).

we are trying to model, then what matters clearly is relative rather than absolute income. In this vein we should note that it is relative rather than absolute inequality that, at least in more developed economies, can be connected to social harm (Wilkinson *et al.* [2011]), thus viewing the models in terms of relative income inequality would not necessarily blunt their normative implications. However, such a strategy would involve quite a radical re-interpretation of the models. It would also serve to conflate the absolute and relative income concepts: ideally we would like models of the absolute income from which we can calculate the relative income. We will leave this the question of whether the kinetic exchange models could be interpreted in terms of relative income also open for further analysis.

3.3 Exchange Dynamics

Recall that the statistical mechanics and econophysics models both involve simple particles or agents that have only one property. For the DY model of income this property is money. Despite its simplicity the DY exchange process does share some important features with real exchanges: in a real exchange (ignoring taxes), the amount of *money* one agent has increases exactly as much as the other agent's money decreases. Of course, typically some goods or services travel in the opposite direction to the money, but if we were just focusing on the distribution of money then what matters is that one agent is poorer and one richer. Chakrabarti *et al.* ([2013]) also discuss similar models with commodities, but we don't need to discuss that here. So even though real interactions are very far removed from the interactions of the kinetic exchange model, the dynamics do capture one important feature of real interactions. What is striking is that it appears that this is all that is required to capture at least some of the stylised facts about the aggregate distribution of income. As Yakovenko and Rosser ([2009]) put it:

> [The DY model] only keeps track of money flow, and does not keep track of what goods and services are delivered. One reason for this is that many goods, e.g. food and other supplies, and most services, e.g. getting a haircut disappear after consumption. Because they are not conserved, and also because they are measured in different physical units, it is not very practical to keep track of them. In contrast, money is measured in the same unit...and is conserved in local transactions, so it is straightforward to keep track of money flow. It is also important to realize that an increase in material production does not produce an automatic increase in money supply. The agents can grow apples on trees, but cannot grow money on trees. Only a central bank has the monopoly on changing the monetary base.

For the gas case the single property playing the same role as money is translational kinetic energy, which, since all molecules in the model have identical mass, is given simply by speed. For polyatomic gases at least, a substantial fraction of molecular energy resides in modes other than translational kinetic energy (Poling *et al.* [2001]). For instance, there can be intra-molecular vibrational and rotational kinetic energy related to the internal molecular structure (there can also be energy related to inter-molecular rotations and vibrations). What is more, real collisions in such gases can involve transfer between such internal degrees of freedom and the translational modes, and such transfers play an important role in determining the outcome of a binary collision between molecules (provided the temperature is not too low) (Trusler [2011]). Consider the analogy of the simplified Boltzmann system: the particles in the gas model are frictionless, and they are not spinning, so when they bounce into each other, the *only* change that occurs is changes to the speed (and direction) of travel. If real spheres

(with friction) were to bump into each other, there would be transfers of rotational energy as well: a glancing hit would cause the spheres to spin more than they did before collision.

We thus see that the elastic transfers of translational momentum that are fundamental to the Boltzmann kinetic exchange model are also heavily idealised exchange processes. There are some monatomic gases (for example, noble gases) that can approximately be taken to possess no internal energy communicable via collisions, and for these gases the Boltzmann model does give a good fit to data on many thermodynamic properties (Chapman and Cowling [1991], §12-14). Thus, one could restrict the target systems relevant to the Boltzmann model to be monatomic gases, and take the idealisation that all kinetic energy is exchangeable via linear momentum transfer to be justified as a legitimate approximation to the real collisions. However, one could also take a different route: even if the idealisation is not an approximation to the exchanges in the target (gases in general, including polyatomic gases) this does not invalidate the model. Rather it opens up natural avenues for 'de-idealisations' that take into account more of the features that in fact make a difference to the behaviour of the target system. In other words, we may take the idealisations to be Galilean idealisations whose simplicity is a pragmatic virtue but may ultimately be removed by 'de-idealising' the model in question (see McMullin [1985], Cartwright [1989]). More specifically we can refine the 'smooth hard sphere collision' exchange mechanism to include dynamical features that better take account of relevant features of real collisions. Following, Chapman and Cowling ([1991], §11), for example we can consider a model for polyatomic gases that was proposed soon after Boltzmann's original paper. In the 'Bryan-Pidduck' model (Bryan ([1894]), Pidduck [1922]) we consider 'rough' elastic spherical molecules that possess rotational energy and are such that in a collision the relative velocity of the spheres at their point of impact is exactly reversed. In such models we have increased transport of rotational energy since rough spheres will knock back molecules with which they collide whilst at the same time transferring part of their rotational energy to them. Although not a more realistic representation of collisions in a polyatomic gas, this adaptation of the Boltzmann model leads to an idealisation which more closely approximates real exchanges in the sense that it better approximates the effects of the features that have been idealised away.

With this story in mind, consider the analogue 'exchange idealisation' in the econophysics case. A feature of real economic agents is that typical monetary exchanges do not relate to the entirety of an agent's income. Agents have other assets that are not captured in the randomised exchange process: just like molecules have other forms of kinetic energy not captured in the linear momentum transfer exchange process. Only in extreme examples would a real interaction between two agents involve all the money of the agents. Thus one idealisation involved in the DY model is that the agent-agent 'collisions' allow for such total exchanges of money. Just like in the Boltzmannian case it is difficult to justify this idealization as a legitimate approximation to real interactions without severely limiting the target systems to which we want to apply the model. However, just like in the Boltzmannian case this does not invalidate the model rather it is promising to initially treat the model as containing a Galilean idealization which, then, gives us a natural starting point for developing a new more realistic model (through de-idealisation of the old model). And this is what econophysicists in fact do: one simple strategy for 'de-idealising' the DY-model is to allow agents to 'save' money by only putting a certain proportion of their total money in the pot to be randomly exchanged during 'collisions' (Chakrabarti *et al.* [2013]). Like the rough spheres and polyatomic gases example, this need not be an adjustment that makes the model representationally more accurate. Rather it may produce a model where the refined idealisations better approximate the effects of the features that have been idealised away. In the following section we will see

that such a simple adjustment to the model can have very significant effects on the match between the stable income distribution and real data.

4 Fat Tails and Savings

In the previous section we occupied ourselves with a particular simple model for the emergence of income distributions in economic systems. Here we will consider a slightly more sophisticated variant of the 'collision' model, called the CCM model. In fact, econophysicists have proposed a strategy to de-idealise the assumption that all money is exchangeable (i.e. by allowing savings, see Section 4 below); interestingly, this specific idealisation is already sufficient to alter the shape of the stable income distribution. Most significantly, this new distribution is claimed to match that observed in real data more closely: not only does the distribution match the bulk, but it also gets the fat tail right. The implications of such a 'congruence' with real data for the evaluation of the status of these models will be considered in the following section.

Following Chakrabarti *et al.* ([2013]) we can characterise the observed regularities in real income distributions as follows: if $P(x)dx$ is the probability that a randomly chosen individual within a population is found to have a income between x and $x + dx$, then the functional form of $P(x)$ is:

$$P(x) \approx x^n e^{\frac{-x}{T}} \text{ for } x < x_c \tag{4.1}$$

$$P(x) \approx x^{-\alpha-1} \text{ for } x \geq x_c \tag{4.2}$$

where n and α are two exponents and T is a scaling factor. The distribution for $x < x_c$ is an exponential distribution and has as a special case the Boltzmann-Gibbs distribution familiar from physics of classical gases at equilibrium. This connection between income distributions and the kinetic theory of gasses was of course exemplified by the models considered in the previous section. The distribution for $x \geq x_c$ is also of a well known form. It is a power law or Pareto distribution. Typically, empirical data imply a value for x_c such that around 90% of the population follows the exponential distribution, and around 10% the Pareto distribution. Since the Pareto distribution characteristically drops off more slowly than the exponential distribution, income distributions display a 'fat tail' effect similar to those studied in other areas of econophysics, particularly financial markets.[16]

Given that we accept the evidence in favour of real income distributions displaying the 'stylised fact' of a bulk exponential distribution and a power law tail, is such a distribution recoverable from collision type models like those considered in the previous section? The prima facie surprising answer is, yes. What is more the *only* modification to the DY model required to reproduce the power-law tail is the addition of a saving propensity to the agents. Explicitly, following Chakrabarti *et al.* ([2013]), we introduce a parameter λ_i which corresponds to a 'savings propensity' of the ith agent. Our collision dynamics now follow the adjusted equations:

$$m_i(t + 1) = m_i(t) + \Delta m \tag{4.3}$$

$$m_j(t + 1) = m_j(t) - \Delta m \tag{4.4}$$

where

$$\Delta m = (1 - \lambda_j)\epsilon_{ij}m_j(t) - (1 - \lambda_i)(1 - \epsilon_{ij})m_i(t) \tag{4.5}$$

---

[16]Mitzenmacher ([2003]) points out that power law distributions and lognormal distributions often look very similar, and summarises the history of this thought in a number of disciplines.
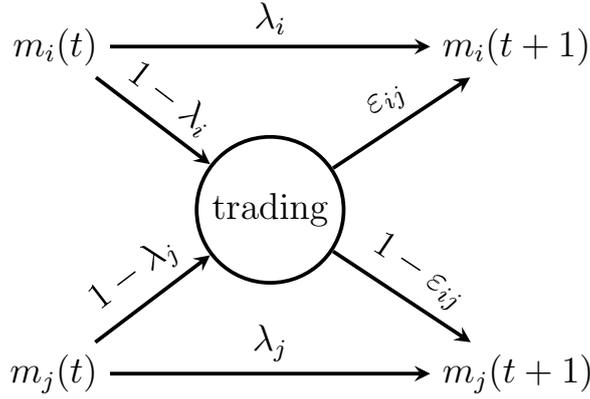
Figure 3: The CCM dynamics (after Chakrabarti et al. [2013], Figure 4.1, p.56)

The effect of the savings propensity is thus to treat the agents as putting some fraction of their money aside before the collision, so that that they each only put a fraction of $1 - \lambda_i$ into the pot to be randomly distributed. Although the equations of this 'CCM' model are only slightly more complicated than the DY model, these complications render an analytical solution much less tractable (although see Chakraborti and Patriarca [2009]). Rather, such models are typically studied using 'Monte Carlo' simulations. Such simulations give a steady state distribution that, for suitably large N and t, and a heterogeneous distribution of λ_i and m_i, reproduces the relevant stylised fact: i.e. bulk exponential and power-law tail.

5 Evaluation

The starting point of our analysis was the question of whether or not kinetic exchange models for income represent (in some sense) examples of an illegitimate or inherently problematic modelling practice. In the previous sections we have analysed two such models in detail, and made what we hope is an instructive comparison with the basic Boltzmannian kinetic model of gases. The broad conclusion of our analysis thus far is that when the same (or very similar) model is applied in different modelling contexts the idealised features that are in common (or very similar) may be sanctioned by very different justificatory arguments. The conservation assumption in the exchange models in an example of this.

A second, rather general, conclusion is that it does not seem particularly plausible to take a model to be invalidated simply by the fact that it involves idealising features with no representational correspondence to 'reality'. Call 'naïve representationalism' the view that holds with respect to a model m if and only if all features of model m, including the idealised ones, can be interpreted as bearing a correspondence (or, at least an approximate correspondence) to features of the target system. Within naïve representationalism the only legitimate idealisations are approximations, i.e. approximate correspondences between model and features of the target system. With good reason, such a restrictive view on models and idealisation is not particularly popular within the philosophical literature; although see Norton ([2012]) for related arguments. As we saw above, even in the case of the gas model we have idealising assumptions that cannot be straightforwardly interpreted as approximately true. Clearly, under naïve representationalism the econophysics models would be judged as inherently problematic, but so would most of modelling in physical science! Our analysis thus adds further weight to the consensus that a more flexible view of idealisation in science is well justified.

At the other end of the spectrum from naïve representationalism, one might consider a maximally permissive view of idealisations in which it is not required that *any* features of the model bear an approximate correspondence to features of the target system. Rather, we consider models such as DY and CCM to be 'generative models' in the sense of Frank ([2009]):

> generative models can be thought of as tools by which we discover important maximum entropy attractor distributions. Once we have found such distributions by a generative model, we may extract the informational constraints that define the pattern[...]The original generative model no longer has special status – our greatest insight resides with the informational constraints that define the maximum entropy distribution.

Given the importance of entropy maximisation to analytical understanding of the stable distributions of both income models, this view might seem particularly attractive. It is, however, a rather extreme position that amounts to a form of rather sparse instrumentalism that is reminiscent, with some irony, of Milton Friedman's economic positivism (Friedman [1953]). In fact, one can see Frank (and arguably Friedman) as committing something akin to what Stein ([1989]) calls 'the fallacy of nothing but', only in the context of social rather than physical science. Just as we should reject a 'trite instrumentalism' that takes a physical theory as 'nothing but an instrument for calculating the outcomes of experiments', we should not accept that social science models are 'nothing but' tools for discovering the constraints that define maximum entropy distributions. Moreover, since we are dealing with models of inequality, which is (presumably) something we would like to diminish, then the conceptual shallowness of 'trite instrumentalism' is particularly worrying. To be of any great use in practice, social science models must give us resources for guiding policy interventions and thus involve some form of causal or counterfactual information. In this spirit, we would follow the insistence of Brock ([1999]) that 'the underlying econometric identification problem of uncovering and estimating the underlying causal data generating mechanism must be dealt with. While scaling law type regularities give us clues, much more must be done to learn the underlying structure.'

This brings us to two 'intermediate' views of idealisations that are of particular relevance to our analysis: the 'minimalist' and 'dispositionalist' conceptions of models. According to minimalism, idealised models truthfully represent (i) a minimal set of explanatorily relevant factors and (ii) idealised assumptions signal that some factor is not explanatorily relevant (Strevens ([2008]), Weisberg ([2013])). The idealised model is taken to represent all, and only, the explanatorily relevant factors for the target phenomenon, and the model ignores, in the form of idealised assumptions, irrelevant factors. As Strevens ([2008]) argues at length, the idealisations in the gas model can plausibly be interpreted in accord with the minimalist interpretation: for the gas, it does seem to be reasonable to take the idealisations discussed in Section 3 as statements about what does not make a difference for the occurrence of the target phenomenon. However, the same does not seem to hold for the income model. As we pointed out above, it is far from clear that real economic features such as many-agent interactions and growth in the money supply are 'explanatorily irrelevant' to the distribution of incomes. Thus, we do not have good reasons to believe that Strevens' version of minimalism in fact applies to exchange models of income. This is, of course, not a knock-down against minimalism. However, the burden of argument is on the minimalists to show that, for instance, many agent interactions are in fact irrelevant to the core mechanism that leads to the distribution of income we observe.

Let us turn to dispositionalism. According to dispositionalism, idealised models truthfully represent the disposition of a (physical or economic) system to behave, if other disturbing causes were absent (Cartwright ([1989]), Hüttemann ([2004])). An idealised assumption describes an often counterfactual situation in which a particular factor is absent and the target system is isolated from the influence of that particular factor. A dispositionalist account would assert that the DY-model describes a disposition to behave in the absence of, for instance, many-agent interactions. Unlike the minimalist, the dispositionalist is not committed to the claim that many-agent interactions are explanatorily irrelevant. However, the dispositionalist faces another problem: they have to justify how the income model is applicable in 'non-ideal' situations, i.e. in the actually quite frequent kind of situation in which many-agent interactions in fact occur (Earman and Roberts [1999]). Meeting this challenge is more difficult in the case of the income model than in the case of the gas model, because, unlike in the gas model, there are no general dynamical laws in the income model that might help us to determine what will happen if 'disturbing factors' are in fact present and, thereby, guide the application of the model in a 'non-ideal' situation (see Hüttemann [2014] for an elaboration of this point).

Given that we have rejected an account of the income models in terms of 'trite instrumentalism', and found accommodation of the models within both the minimalist and dispositionalist frameworks problematic, one might conclude that the 'Worrying Trends' criticism of these models has been justified. However, we think this is a little too fast. Although we saw good reason to reject a 'trite' form of instrumentalism, we have not ruled out a richer form of instrumentalism, powerful enough to treat models as tools for motivating interventions. Moreover, if we accept that highly idealised 'toy models' (such as Schelling's model of segregation, the Lotka-Volterra predator-prey model or the 'MIT' bag model of quantum chromodynamics) can give 'how-possibly' explanations of the relevant target phenomenon, then such a 'toy model' account might be applicable to the DY and CCM models also.[17] Econophysics is, in the end, an example of new and innovative cross-disciplinary field, involving the application of paradigms and tools borrowed from theoretical and statistical physics within the social sciences. We should not, therefore, be entirely surprised that at least some of the models of econophysics do not always comfortably fit within existing accounts of models and idealisations in the physical and social sciences: this might be taken to be as much of a problem for the existing accounts as it is for the econophysics models. The econophysics of social phenomena offers new territory for philosophers of science to explore, and the current paper in intended principally as a first step[18] in the philosophical analysis of this exciting, yet perplexing, new enterprise.

Acknowledgement

---

[17] See Hartmann ([1998]), Grüne-Yanoff ([2009]).

[18] We again point to the existence of admirable work of Kuhlmann [2006, 2014], Rickles [2007, 2008] and Casini ([2014]) focusing upon financial econophysics.

References

Atkinson, A. B. and Bourguignon, F. [2000]: *Handbook of Income Distribution*, vol. 1 Elsevier.

Banerjee, A. and Yakovenko, V. M. [2010]: 'Universal patterns of inequality', *New Journal of Physics*, **12**(7), pp. 075032.

Bertola, G., Foellmi, R. and Zweimüller, J. [2014]: *Income Distribution in Macroeconomic Models*, Princeton University Press.

Boltzmann, L. [1970]: *Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen*, Springer.

Brock, W. [1999]: 'Scaling Laws in Economics: A Reader's Guide', *Industrial and Corporate Chance*, **8**(3), pp. 409–446.

Brown, J. [1976]: 'The Mathematical and Statistical Theory of Income Distribution', in A. B. Atkinson (*ed.*), *The Personal Distribution of Income*, London: George Allen & Unwin LTD.

Bryan, G. H. [1894]: 'On the application of the determinantal relation to the kinetic theory of polyatomic gases', *Rep. British Assoc. Adv. Sci*, **64**, pp. 102–106.

Cartwright, N. [1989]: *Nature's Capacities and Their Measurement*, Oxford: Oxford University Press.

Casini, L. [2014]: 'Not-so-minimal Models. Between Isolation and Imagination', *Philosophy of the Social Sciences*, **44**(5), pp. 646–672.

Chakrabarti, B. K., Chakraborti, A., Chakravarty, S. R. and Chatterjee, A. [2013]: *Econophysics of income and wealth distributions*, Cambridge University Press.

Chakraborti, A. and Patriarca, M. [2009]: 'Variational principle for the Pareto power law', *Physical review letters*, **103**(22), pp. 228701.

Champernowne, D. G. [1953]: 'A model of income distribution', *The Economic Journal*, pp. 318–351.

Chapman, S. and Cowling, T. [1991]: *The Mathematical Theory of Non-uniform Gases*, Cambridge University Press.

Christiano, L. J. [1994]: 'Modeling the liquidity effect of a money shock', in *Inventory, Business Cycles and Monetary Transmission*, Springer, pp. 61–124.

Cont, R. [2001]: 'Empirical properties of asset returns: stylized facts and statistical issues', *Quantitative Finance*, **1**(2), pp. 223–236.

Drăgulescu, A. and Yakovenko, V. M. [2000]: 'Statistical mechanics of money', *The European Physical Journal B-Condensed Matter and Complex Systems*, **17**(4), pp. 723–729.

Earman, J. and Roberts, J. [1999]: 'Ceteris paribus, there is no problem of provisos', *Synthese*, **118**(3), pp. 439–478.

Emch, G. G. and Liu, C. [2002]: *The Logic of Thermostatistical Physics*, Springer.

Frank, S. A. [2009]: 'The common patterns of nature', *Journal of evolutionary biology*, **22**(8), pp. 1563–1585.

Friedman, M. [1953]: 'The methodology of positive economics', in *Essays in Positive Economics*, University of Chicago Press.

Gabaix, X. [1999]: 'Zipf's law for cities: an explanation', *Quarterly journal of Economics*, pp. 739–767.

Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H. E. [2005]: 'Institutional investors and stock market volatility', Tech. rep., National Bureau of Economic Research.

Gallegati, M., Keen, S., Lux, T. and Ormerod, P. [2006]: 'Worrying trends in econophysics', *Physica A: Statistical Mechanics and its Applications*, **370**(1), pp. 1–6.

Gingras, Y. and Schinckus, C. [2012]: 'The institutionalization of econophysics in the shadow of physics', *Journal of the History of Economic Thought*, **34**(1), pp. 109–130.

Grüne-Yanoff, T. [2009]: 'Learning from minimal economic models', *Erkenntnis*, **70**(1), pp. 81–99.

Hartmann, S. [1998]: 'Idealization in quantum field theory', in N. Shanks (*ed.*), *Idealization in Contemporary Physics*, vol. 63, Amsterdam: Rodopi, pp. 99–122.

Hogan, J. [2005]: 'Why it is hard to share the wealth', *New Scientist*.

Homburg, S. [2015]: 'Critical remarks on Piketty's Capital in the Twenty-first Century', *Applied Economics*, **47**(14), pp. 1401–1406.
<http://dx.doi.org/10.1080/00036846.2014.997927>

Hüttemann, A. [2004]: *What's Wrong with Microphysicalism?*, Routledge.

Hüttemann, A. [2014]: 'Ceteris paribus laws in physics', *Erkenntnis*, **79**(10), pp. 1715–1728.

IMF [2014]: 'Money and quasi money growth', Tech. rep., International Monetary Fund.
<http://data.worldbank.org/indicator/FM.LBL.MQMY.ZG/countries?page=1>

Jaynes, E. T. [2003]: *Probability Theory: The Logic of Science*, Cambridge university press.

Jovanovic, F. and Schinckus, C. [2013a]: 'The emergence of econophysics: A new approach in modern financial theory', *History of Political Economy*, **45**(3), pp. 443–474.

Jovanovic, F. and Schinckus, C. [2013b]: 'Towards a transdisciplinary econophysics', *Journal of Economic Methodology*, **20**(2), pp. 164–183.

Kesten, H. [1973]: 'Random difference equations and renewal theory for products of random matrices', *Acta Mathematica*, **131**(1), pp. 207–248.

Kuhlmann, M. [2006]: 'How Do Microscopic Models of Financial Markets Explain?', *http://philsci-archive.pitt.edu/2788/*.
<http://philsci-archive.pitt.edu/2788/>

Kuhlmann, M. [2014]: 'Explaining Financial Markets in Terms of Complex Systems', *Philosophy of Science*, **81**(1), pp. 1117–1130.

Ladley, D. [2013]: 'Zero intelligence in economics and finance', *The Knowledge Enginerring Review*, **27**(2), pp. 273–286.

Ladyman, J. [2013]: 'Idealization', in M. Curd and S. Psillos (*eds*), *The Routledge companion to philosophy of science*, Routledge, chap. 33, pp. pp.356–66.

Lux, T. [2005]: 'Emergent statistical wealth distributions in simple monetary exchange models: a critical review', in *Econophysics of Wealth Distributions*, Springer, pp. 51–60.

Lyon, A. [2011]: 'Deterministic probability: neither chance nor credence', *Synthese*, **182**(3), pp. 413–432.
<http://dx.doi.org/10.1007/s11229-010-9750-2>

Lyon, A. [2013]: 'Why are Normal Distributions Normal?', *The British Journal for the Philosophy of Science*, **65**(3), pp. 621–649.

Mandelbrot, B. [1960]: 'The Pareto-Levy law and the distribution of income', *International Economic Review*, **1**(2), pp. 79–106.

McCauley, J. L. [2006]: 'Response to "worrying trends in econophysics"', *Physica A: Statistical Mechanics and its Applications*, **371**(2), pp. 601–609.

McMullin, E. [1985]: 'Galilean idealization', *Studies in the History and Philosophy of Science*, **16**(3), pp. 247–273.

Mincer, J. [1958]: 'Investment in human capital and personal income distribution', *The Journal of Political Economy*, pp. 281–302.

Mirowski, P. [1992]: *More heat than light*, Cambridge University Press.

Mitzenmacher, M. [2003]: 'A Brief History of Generative Models for Power Law and Lognormal Distributions', *Internet Mathematics*, **1**(2), pp. 226–251.

Norton, J. D. [2012]: 'Approximation and Idealization: Why the Difference Matters', *Philosophy of Science*, **79**(2), pp. 207–232.

Pidduck, F. [1922]: 'The kinetic theory of a special type of rigid molecule', *Proceedings of the Royal Society of London. Series A*, **101**(708), pp. 101–112.

Piketty, T. [2014]: *Capital in the 21st Century*, Harvard University Press.

Poling, B. E., Prausnitz, J. M., John Paul, O. and Reid, R. C. [2001]: *The Properties of Gases and Liquids*, vol. 5 McGraw-Hill New York.

Rickles, D. [2007]: 'Econophysics for philosophers', *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **38**(4), pp. 948–978.

Rickles, D. [2008]: 'Econophysics and the complexity of financial markets', in *Handbook of the Philosophy of Science: Philosophy of Complex System*, vol. 10, North Holland: Elsevier, pp. 133–152.

Rosser, J. B. [2010]: 'Is a transdisciplinary perspective on economic complexity possible?', *Journal of Economic Behavior & Organization*, **75**(1), pp. 3–11.

Sinha, S. and Chakrabarti, B. [2012]: 'Econophysics: an emerging discipline', *Econ Polit Wkly*, **46**, pp. 44.

Stanley, H. E. and Mantegna, R. N. [2000]: *An introduction to econophysics*, Cambridge University Press, Cambridge.

Stein, H. [1989]: 'Yes, but… Some Skeptical Remarks on Realism and Anti-Realism', *Dialectica*, **43**(1-2), pp. 47–65.

Strevens, M. [2008]: *Depth: An Account of Scientific Explanation*, Harvard University Press.

Tinbergen, J. [1975]: *Income Distribution: Analysis and Policies*, Elsevier.

Trusler, J. [2011]: 'Kinetic Theory of Gases', *http://www.thermopedia.com/content/907/*. <http://www.thermopedia.com/content/907/>

Uffink, J. [2014]: 'Boltzmann's Work in Statistical Physics', in E. N. Zalta (*ed.*), *The Stanford Encyclopedia of Philosophy*, fall 2014 edition.

von Neumann, J. and Morgenstern, O. [1944]: *Theory of Games and Economic Behavior*, Princeton University Press.

Walsh, C. E. [2010]: *Monetary Theory and Policy*, MIT press.

Weatherall, J. O. [2013]: *The Physics of Wall Street: A Brief History of Predicting the Unpredictable*, Houghton Mifflin Harcourt.

Weisberg, M. [2013]: *Simulation and Similarity: Using Models to Understand the World*, Oxford University Press New York.

Wiens, B. L. [1999]: 'When Log-Normal and Gamma Models Give Different Results: A Case Study', *The American Statistician*, **53**(2), pp. 89–93.

Wilkinson, R. G., Pickett, K. and Chafer, C. [2011]: *The Spirit Level*, Tantor Media, Incorporated.

Winsberg, E. [1999]: 'Sanctioning models: The epistemology of simulation', *Science in context*, **12**(02), pp. 275–292.

Wold, H. O. and Whittle, P. [1957]: 'A model explaining the Pareto distribution of wealth', *Econometrica, Journal of the Econometric Society*, pp. 591–595.

Yakovenko, V. M. and Rosser, J. B. [2009]: 'Colloquim: Statistical Mechanics of Money, Wealth and Income', *Reviews of Modern Physics*, **81**, pp. 1703–1725.