Dynamics of inequality

Publication of the Panama Papers helped shine a light on the hidden wealth of some of the world's most powerful individuals, putting rising inequality back on the political agenda. But we should not be surprised when inequality increases, say **Alexander Adamou** and **Ole Peters**. Meanwhile, on page 36, **Danny Dorling** offers a response

Rising inequality in many parts of the world is recognised as one of today's greatest social challenges. Oxfam recently reported that 62 people have the same combined wealth as the poorest half of the world's population.[1] That a wealthy elite has benefited disproportionately from the economy in recent years is difficult to dispute.

Often these phenomena are explained in terms of special circumstances. Oxfam draws attention to tax havens, the widespread use of which was exposed by the recent hacking of a Panamanian law firm. Piketty points to a "rentier" class whose capital grows at a privileged rate.[2] We do not dispute these factors, but we point out that inequality may rise even without them. If this is true then policies aimed at limiting or reducing inequality have to intervene actively rather than merely eliminate privilege. Indeed, the tendency for inequality to increase to detrimental levels was recognised by ancient societies. The Bible (Leviticus 25:8–13) speaks of the jubilee year of debt forgiveness, a radical reduction of inequality, every 49 years. Whether a society seeks to control inequality is a political choice. This article offers no political counsel. Instead it investigates how the way we conceptualise change in economics influences our response to the choice, once made.

We explore a simple economic model – standard in mathematical finance – with an inherently unstable wealth distribution. This motivates a fundamental definition of inequality (see equation (4), page 34), which coincides with a known inequality measure, the mean logarithmic deviation (MLD), when wealth reproduces multiplicatively. The MLD was previously derived without considering the economic mechanisms by which inequality arises. Our new derivation gives it a new interpretation, identifying it as the natural inequality measure for an economy predicated on compound growth of capital investments.

Alexander Adamou and **Ole Peters** are fellows of the London Mathematical Laboratory. Peters is also a professor of the Santa Fe Institute

Unsolved problems

Despite the obvious importance of how societies share what they produce, our scientific understanding of economic inequality is primitive. Two basic problems remain unsolved.

Firstly, there is no consensus about how to measure it. Perhaps this is not as great a failure as it sounds. Inequality is a nebulous concept. It tracks the broadness of the distribution of the economic resources available to an individual, usually measured by wealth or income. (By "distribution" we mean a probability density function, rather than the act of sharing of a resource among a population.) We know, for example, that inequality is highest when one person has everything and lowest when everyone has the same, and that it increases when resources are transferred from poorer to richer. It is measured by condensing the distribution into a single number – the inequality metric – whose properties match this intuition.

As there is no precise definition of inequality, there is no unique choice of metric. Using data from the US Census Bureau, Figure 1 plots four common inequality metrics for US household income from 1967 to 2013.[3] All four metrics tell a similar story, that inequality grew consistently. This similarity is unsurprising: they are, after all, derived from the same underlying distributions. However, differences between the metrics are visible. They correspond to subtly different interpretations of the inequality concept.

Secondly, comparatively little attention has been given to how economic distributions, and therefore inequality, change over time. Figure 1 tells us that the US income distribution *has* changed over time because all the inequality metrics derived from it have increased. Regardless of whether we consider this desirable, it would surely be useful to understand what causes such changes and what actions society might take if it wanted to control them.

Main image: fredcardoso/iStock/Thinkstock

This is a specific example of a more general question, which goes to the root of how economic science is conducted: how do changes in the economy occur? There are two different worldviews, which depend on whether we believe the economy is in equilibrium and which lead to different types of explanations for observed phenomena.

Equilibrium

Many different notions of equilibrium exist in economics, but most people first encounter the basic concept in a physics or chemistry class. Our definition will be drawn from these fields. Measuring the wealth inequality of a population is rather like measuring the temperature of a gas in a container. In both systems the observable is determined by the distribution of some microscopic variable: temperature by the distribution of molecular kinetic energy; wealth inequality by the distribution of personal wealth. We are interested in a small number of macroscopic observables, which summarise the whole system. Temperature and inequality are examples of these.

Some physical systems are well described by mathematical models that possess a stable distribution. We call these "equilibrium systems". Empirical distributions will resemble the stable model distribution if we let the system come to rest. All macroscopic observables will similarly stabilise. Once this has happened, we say the system is "in equilibrium".

We speak of "non-equilibrium systems" when these concepts do not apply. This may be because it will take longer for the system to come to rest than we will be observing it, or because our mathematical model does not predict a stable distribution in the first place. In this case macroscopic observables can change when measured over time.

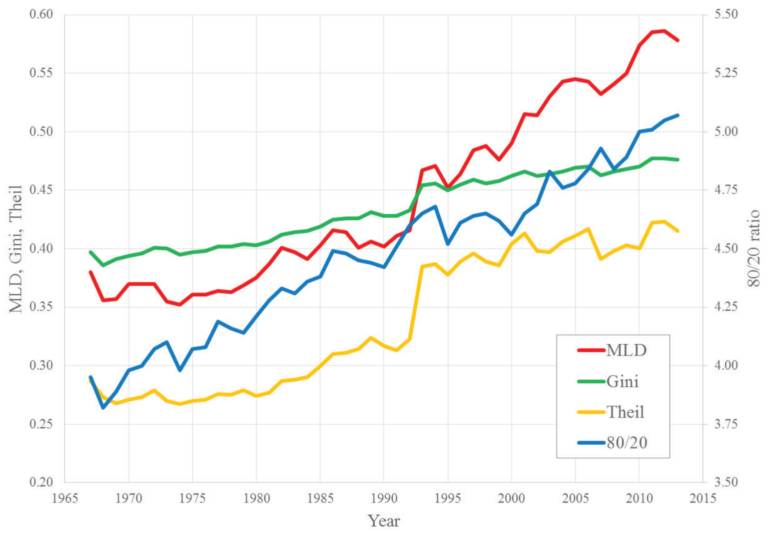

FIGURE 1 Evolution of four relative metrics of US household income inequality: mean logarithmic deviation (MLD, red); Gini coefficient (green); Theil index (yellow); 80/20 ratio (blue). Data from US Census Bureau[3]

Controlling inequality is a political choice, but how we conceptualise change in economies influences our response to this choice

The equilibrium worldview

In one worldview economies are usually in equilibrium. The stable model distributions are determined by prevailing economic conditions. In the case of wealth and income models, parameters for taxation and welfare, employment, housing, investments, and so on may have to be included. We will not worry about the details. What matters is that these parameters can change, causing the stable distribution to change.

The equilibrium worldview holds that, when this happens, the observed distribution converges to the new stable distribution much faster than the parameters change. We do not observe the brief transitions from one stable distribution to the next. Instead, when we observe inequality, it remains static while the distribution remains static. The only thing that can change it is a shift to another distribution.

The implication is that inequality is, by default, a stable observable. It can change only in response to economic conditions. If observations show changes, as in Figure 1, they require explanation in those terms. This is the dominant approach in mainstream economics. It leads to a certain type of answer. For example, to the question of why inequality is rising, we have the answers that the rich are paying less tax and their capital is growing faster than before. These are possible answers, but they might lead us to imagine that inequality would stabilise if we got rid of tax havens and privileged investments. This may not be true.

The non-equilibrium worldview

The alternative worldview is that economies are non-equilibrium systems. Empirical distributions never stabilise and macroscopic observables can change over time. Here the non-equilibrium behaviour cannot be ignored. We can include a backdrop of slowly changing economic conditions, represented by slowly varying parameters in a model, but this does not alter the fundamental point. In the non-equilibrium worldview, distributions – and, therefore, observables such as inequality – can change even when conditions do not. Change is a fact of life.

This leads to a different type of science. To the question why inequality is growing (or, more generally, changing) it is now admissible to answer that this is the default behaviour of the system. This does not mean that changes in conditions have no explanatory power, only that they do not have the monopoly on it. Observed trends, such as in Figure 1, can occur in an economy with static conditions.

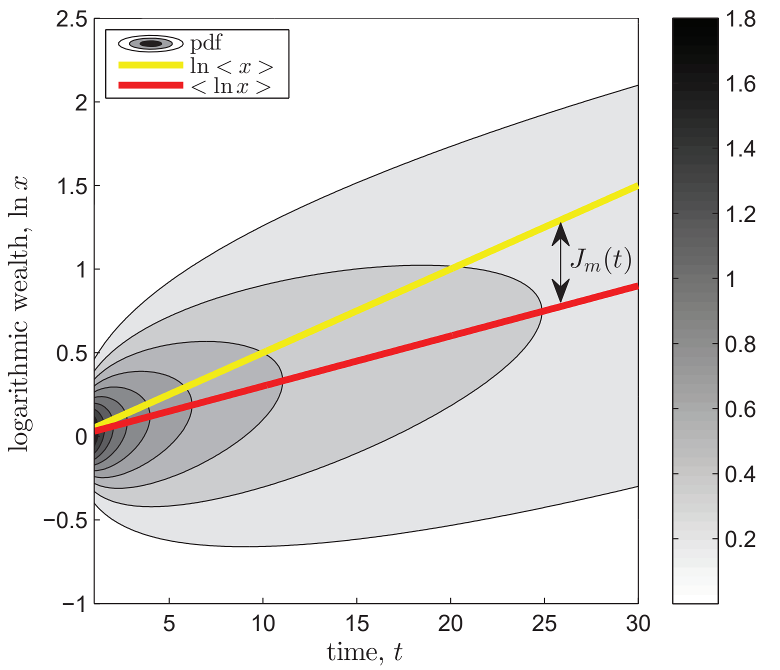

FIGURE 2 Contour plot of the evolution of the probability density function of a GBM with $\mu = 0.05$ per unit time and $\sigma = 0.2$ per $\sqrt{\text{unit time}}$. The evolution of the logarithms of the average (yellow) and typical (red) wealth are also shown. Inequality, J_m, grows as the difference between the two

This has important implications for policy. If we believe our economy is a non-equilibrium system, then we should expect inequality to change. Controlling it requires intervention to alter the nature of the system, for example by systematically reallocating resources.

A simple model

We illustrate the non-equilibrium approach by exploring a toy model. Imagine a population of *N* individuals whose wealths evolve according to some process. For simplicity we will ignore births, deaths and other events that might cause *N* to change in time. If we think about most real-world investments, we see two basic ingredients: the profit or loss is related to the amount of money invested; and it is uncertain. These two ingredients define a general class of processes. One of the simplest is noisy multiplicative growth, modelled as geometric Brownian motion (GBM). Each individual's wealth is governed by the stochastic differential equation

$$dx = x\left(\mu\, dt + \sigma\, dW\right) \qquad (1)$$

where x is the wealth, dx is the change in wealth over time period dt, dW is a normal random variable of mean zero and variance dt, and μ and σ are parameters known as the drift and volatility.

Put simply, wealth changes by a random proportion of itself over each time step. This model obviously lacks detailed realism. It contains none of the reallocation that we know occurs in real societies, such as taxation and public spending. Alone it cannot explain decreases in inequality. Although we can enhance the model to include such effects,[4] it is nonetheless the most-used model for wealth processes in mathematical finance and does reflect some real aspects of wealth evolution. We use it here because it is simple and shares with more realistic models universal properties we want to highlight.

A lot is known about GBM. The wealth of each individual is log-normally distributed, which means that the logarithm of the wealth is normal:

$$\ln x(t) \sim N\left(\ln x(0) + \left(\mu - \frac{\sigma^2}{2}\right)t,\ \sigma^2 t\right) \qquad (2)$$

The variance of $\ln x(t)$ grows over time, so in this model the wealth distribution never stabilises. Instead it broadens continually and any inequality metric derived from it increases continually, much like the actual trends in Figure 1. Furthermore, in this model almost all the wealth eventually belongs to almost none of the people,[5] reminding us of Oxfam's findings.

Growth rates and inequality

In GBM each individual's wealth grows exponentially. Randomness means that some trajectories grow faster than others over short times, but in the model they all converge to a common growth rate, g_{typ}, over time. This is the typical growth rate observed in the population if we wait long enough.

We can also consider the mean wealth,

$$\left\langle x(t)\right\rangle_N \equiv \sum_{i=1}^{N} x_i(t)/N \qquad (3)$$

where $x_i(t)$ is the wealth of the *i*th individual at time *t*. An important feature of GBM is that the mean wealth tends to grow at a faster rate than the typical wealth, $g_{ave} > g_{typ}$, for realistic population sizes, timescales, and parameters.

The existence of two different behaviours – one for a typical member, revealed over time, and another for the aggregate population – is a feature of a general class of models. We call this feature "non-ergodicity".[6] It inspires the following insight: inequality changes when the growth rates of the average and typical wealths are different. If the average wealth grows faster than the typical wealth, then there must be a handful of unusually wealthy individuals growing unusually quickly. This means inequality is increasing. If the converse is true, then inequality is decreasing.

Instead of picking and choosing between the inequality metrics that have been proposed, we can simply *define* inequality as the quantity, *J*, which grows at the difference between these two rates,

$$\frac{\Delta J}{\Delta t} = g_{ave} - g_{typ}, \qquad (4)$$

where ΔJ is the change in *J* over time period Δt. This is a fundamental definition with a clear dynamical interpretation. Although motivated by a property of GBM, it can be applied to other models of wealth or income. The only requirement is that the two growth rates exist.

Re-deriving an old metric

For GBM we can write the two growth rates as

$$g_{ave} = \frac{\Delta \ln\langle x(t)\rangle_N}{\Delta t}, \quad g_{typ} = \frac{\Delta\langle \ln x(t)\rangle_N}{\Delta t} \tag{5}$$

Substituting into the definition of J gives us an expression for the inequality metric in our specific model. We label it J_m because it is the appropriate metric for any multiplicative process. Up to a constant term, it is

$$J_m(t) = \ln\langle x(t)\rangle_N - \langle \ln x(t)\rangle_N \tag{6}$$

that is, the difference between the logarithm of the mean wealth and the mean of the logarithm of wealth. This is precisely the definition of MLD, one of the metrics plotted in Figure 1. Theil proposed it in 1967 as a measure of income inequality based on information theory.[7] He obtained it from a class of measures of the entropy of a data set and even remarked that its choice was "against intuition". We obtain it instead from economic considerations: the dynamics of wealth or income. It is the natural inequality measure under multiplicative growth.

Figure 2 illustrates this graphically in the large N limit. The contour plot shows the probability density function for wealth in the GBM model. This broadens over time, since it is a non-equilibrium system. The yellow and red lines show the logarithms of the average and typical wealths. These grow linearly in time, the former faster than the latter. The difference, which is J_m, therefore also grows linearly:

$$J_m(t) = J_m(0) + \frac{\sigma^2 t}{2} \tag{7}$$

A more meaningful measure

Our approach asks what causes inequality to change. We are familiar with explanations of rising inequality based on differential treatment of members of society. We are also familiar with attempts to decrease inequality by social interventions, such as taxation and public spending. However, we should be mindful that inequality can rise purely as the result of random multiplicative growth in an economy in which all individuals are subject to the same, static conditions. We can view this as an underlying tendency, which most societies – going back to biblical times – have sought actively to counter. This view is impossible under equilibrium assumptions.

Observations can be interpreted in a more nuanced way. If we view inequality metrics simply as summaries of distributions, then the interpretation of Figure 1 is nothing more than the crude observation that income inequality in America increased. This is not without use – societies should monitor their internal trends – but we can say more. If income is well modelled as a multiplicative process, then equation (7) tells us that we can expect to see J_m increase linearly in time. Data showing a different trend suggest other effects, such as reallocation, at work.

Finally, we express quantitatively the interplay between inequality and the two fundamental growth rates in a population: those of the average and typical wealth or income. For income, many countries report time series of mean and median values, whose growth rates can be computed and will resemble g_{ave} and g_{typ}. Equation (4) shows how these are related dynamically to J. This could help policy-makers craft rational and effective reallocative policies.

For example, a government may wish to hold g_{typ} above some level ($g_{typ} > g$) during a recession (negative g_{ave}) in order to insulate typical members of the population from hardship. Equation (4) says it can achieve this by reallocating at a minimum rate:

$$-\frac{\Delta J}{\Delta t} > g - g_{ave} \tag{8}$$

That J has a lower bound (at perfect equality) defines the limits of this strategy. It can be effective only while there exists in the economy capacity to redistribute. If J approaches its minimum such that equation (8) cannot be sustained, then policy-makers must look elsewhere. Conversely, if J is far from its minimum, then the scope for intervention is great. Such analyses should, at the very least, complement the customary ideological debates.

Conclusion

The way we conceptualise change is central to how we conduct economic science. Assuming equilibrium places a strong constraint on our models and explanations of observed phenomena. Lifting this constraint puts more tools at our disposal. By exploring a toy model of wealth evolution, we showed how the non-equilibrium worldview helps us understand changes in distributions. It motivates a fundamental definition of inequality, breathes new meaning into an established measure, and enhances our interpretation of data. It could even provide a quantitative basis for more effective policy-making. ■

References

1. Hardoon, D., Fuentes-Nieva, R. and Ayele, S. (2016) *An Economy for the 1%: How Privilege and Power in the Economy Drive Extreme Inequality and How This Can Be Stopped*. Oxford: Oxfam International. bit.ly/1n6OfHv

2. Piketty, T. (2014) *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.

3. DeNavas-Walt, C. and Proctor, B. D. (2014) *Income and Poverty in the United States: 2013*, United States Census Bureau Current Population Reports P60-249. Washington, DC: US Government Printing Office.

4. Berman, Y., Peters, O. and Adamou, A. (2016) *Far from equilibrium: Wealth reallocation in the United States*. arXiv:1605.05631

5. Bouchaud, J.-P. and Mézard, M. (2000) Wealth condensation in a simple model of the economy. *Physica A*, **282**, 536–545.

6. Peters, O. and Gell-Mann, M. (2016) Evaluating gambles using dynamics. *Chaos*, **26**, 023103.

7. Theil, H. (1967) *Economics and Information Theory*. Amsterdam: North-Holland.