Wealth Dynamics: Reallocation, Debt and Bankruptcy

Abi Coleman[a], **James Price**[a], **and Jake Thomas**[a]

[a]Mathematics for Real-World Systems II, University of Warwick, CV4 7AL Coventry, United Kingdom

This manuscript was compiled on June 15, 2020

Wealth inequality is growing globally, but few models are capable of representing this fully. As such, we would like to be able to gain insight into the processes causing this growth, rather than attempting to create a fully predictive model. In this paper we examine several simple conceptual models which capture some of the dynamics of wealth inequality, beginning with multiplicative random walks (MRWs). We explore the relationship between the discrete (MRWs) and continuous cases (Geometric Brownian Motion (GBM)), in order to simulate multiple wealth trajectories. We introduce reallocation using a model first investigated in Marsili et al.'s 1998 paper (1). Reallocation can be thought of as taxes when it is positive (money flows from rich to poor) and rent when negative (money flows from poor to rich). This model represents a collective of GBMs that share their wealth, and we investigate its analytical and numerical properties. We find several novel behaviours in the negative reallocation regime. Finally, we extend this model to include bankruptcy, in order to prevent one unrealistic consequence of the original model where an individual's wealth can tend to negative infinity, and find that the effects are not stabilising as we would expect.

wealth inequality | reallocation | stochastic process | bankruptcy

1. Introduction

Wealth inequality is a global problem; poverty is still prevalent in the world's richest countries. In the UK, 22% of people live in relative low income in 2018/19 (2), meaning they are living with 60% or less of the median income. Furthermore, the Gini coefficient, a measure of inequality which is 0 if everyone is equal and 1 if one person has control of all of the wealth, has increased from 0.272 in 1977 to 0.332 in 2012/13 in the UK (3). Therefore, there is interest in exploring how the distribution of wealth changes over time.

There are many different models which may represent the dynamics of individuals' wealth trajectories, the most basic of which is as a multiplicative random walk (MRW) with drift. (4) In contrast to a simple random walk, this model allows wealth to grow and decline by percentages rather than additive increments; this is more accurate to the real world. (5)

The continuous version of an MRW can be modelled as a geometric Brownian motion (GBM), which is explored in the first section. This includes a noise term which accounts for fluctuations in real world wealth, such as having to replace a boiler, a random one time expense.

GBM captures some features of real wealth trajectories. In GBM the mean is controlled by the individuals with the greatest wealth. This can be orders of magnitude higher than the median value, and is similar to the real world. Furthermore, wealth data in the real world has a distribution with heavy tails, a property which GBM satisfies. Therefore, GBM is an appropriate model to begin from, and is the first model we will investigate analytically in this paper. The relationship between GBM and MRWs will be shown, and will be used to inform our understanding of GBM as a base model and the considerations we must make when simulating multiplicative processes.

However, GBM does not include interaction between individuals; in order to better represent the real world, we include an additional term for reallocation. This term allocates wealth to

Significance Statement

Wealth Inequality is a growing issue globally. Between 1980 and 2016, the poorest 50% of humanity only captured 12 cents in every dollar of global income growth. By contrast, the top 1% captured 27 cents of every dollar. (6) Meanwhile poverty has a devastating impact around the world. (7) Despite this, mathematical models of population level wealth are still in their early stages. Here, we explore a model based on geometric Brownian motion and a simple reallocation mechanism between agents,through which we gain understanding of the intrinsic growth of inequality and how simple mechanisms of interaction, such as tax, can affect this growth. A key benefit of the model is the reallocation parameter: it both quantifies inequality, analogously to classic measures such as the Gini coefficient, and can be directly manipulated within the model for analysis.

each individual depending on how far they are from the mean. With positive reallocation, the mean is attracting, but with negative reallocation, the mean repels trajectories. This model has been used in previous papers (1) (8) such as Berman et al. (9), and was found to be able to reproduce wealth data without overfitting with a reallocation parameter timeseries calculated from the data. Berman et al. (9) found that the reallocation parameter (τ) has been negative in the US since the 1980s, which suggests that on average the overall structure of society is taking wealth from the poor and giving it to the rich. This highlights the importance of the negative reallocation regime. However, prior to this paper there has been little formal investigation into its properties.

We will examine the properties of this model both analytically and empirically, before attempting to improve upon this model by including a concept of bankruptcy.

One feature of GBM with reallocation is that with a negative reallocation parameter, it allows for an individual's wealth to tend to negative infinity. As this does not occur in the real world, we explore whether the model can be improved by including a threshold beyond which an individual's wealth is returned to 1. The investigations occur primarily through simulations, as the added complexity makes the process analytically unwieldy.

2. Multiplicative Random Processes and Simulating Geometric Brownian Motion

Wealth data usually exists as time series with fixed time steps. These time steps are often on the order of months and years due to the difficulty of collecting such data. Therefore, it is natural to initially think of wealth as a discrete time multiplicative process such as a MRW. We will explore the distributions of such MRW and the role rare events play. Before linking MRW to their continuous analogue GBM, which forms the basis of the reallocation model, and discussing the fundamental considerations of simulating random multiplicative processes such as GBM.

A. Multiplicative Random Walks do not always follow a log-normal distribution. We consider a binomial multiplicative random walk which is a sequence of positive real numbers z_1 and z_2 of length N which appear independently with probability p and probability $1-p$, respectively. We are interested in the distribution of the N-fold product, X. Our goal is that this MRW will approximate GBM in the continuum limit.

The classical way of approaching such a problem would be via the log-normal form. However, following the arguments of Redner (10) we will see it is not appropriate in this context. Consider the k-th moment of the product which can be approximated by the integral

$$\langle X^k \rangle \simeq \int_{-N}^{N} e^{\ln p(n)} z_1^{nk} z_2^{(N-n)k} dn \qquad [1]$$

Where $p(n) = \binom{N}{n} p^n q^{N-n}$ is the probability that the product takes the form $z_1^n z_2^{N-n}$. Applying the De Moivre-Laplace Theorem (11) we obtain a Gaussian approximation for $p(n)$

$$p_G(n) \simeq \frac{1}{\sqrt{2\pi Npq}} \exp \frac{(\ln X - \ln X_{mp})^2}{2Npq \left(\ln \frac{z_1}{z_2}\right)^2} \qquad [2]$$

where $X_{mp} = z_1^{Np} z_2^{Nq}$ is the most probable value of the product. By using Laplace's method to evaluate the integral we obtain the familiar log-normal distribution.

$$\begin{aligned} \langle P^k \rangle_G &\simeq \frac{1}{\sqrt{2\pi Npq}} \int_{-\infty}^{+\infty} \exp \frac{-(n-Np)^2}{2Npq} z_1^{nk} z_2^{(N-n)k} dn \qquad [3] \\ &= \exp\left[\left(\frac{Npq}{2}\right) k^2 \left(\ln\left(\frac{z_1}{z_2}\right)\right)^2 + N(p-q)\, k \ln \frac{z_1}{z_2} \right. \\ &\quad \left. + \left(\frac{Nk}{2}\right) \ln z_1 z_2 \right]. \end{aligned}$$

However, we expect the moments of the product to behave like $\left(pz_1^k + qz_2^k\right)^N$. The error arises as Laplace's method is equivalent to finding the maximum of the exponent in equation Eq. (3). This maximum is a poor estimate for the maximum when $p(n)$ has not been approximated.

The true maximum is given by

$$x^\star \simeq \frac{\zeta}{1+\zeta}, \quad \zeta = \frac{pz_1^k}{qz_2^k}$$

while the maximum obtained after the Gaussian approximation is given by

$$x^\star = k \ln \frac{z_1}{z_2} pq + p. \qquad [4]$$

For $z_1, z_2 \not\simeq 1$ or $p \not\simeq \frac{1}{2}$ these values will not agree and for large moments the difference will increase (See Table 1).

Despite this, it is possible to define a MRW which has the correct distribution in the limit.

B. Simulating GBM with a MRW. The first step we take is to transform a given MRW to an additive random walk (ARW). Let X_t be the position of a MRW with time step 1, $z_1 = \frac{1}{z_2} = m$ and initial value x_0 at time t. The process can be broken down into individual steps and written as,

$$X_t = x_0 \prod_{i=1}^{t} X_i$$

Table 1. Comparison of the true first moment (mean) to the one obtained by the log-normal method

z_1	p	True	Gaussian
2	0.5	1.2500	1.2715
5	0.5	2.600	3.6515
20	0.5	10.025	88.872
5	0.55	2.8400	4.9733
5	0.65	3.3200	8.5356
5	0.75	3.800	13.208

True mean is calculated directly. Gaussian mean calculated by equation 3. z_2 is taken as $\frac{1}{z_1}$.

where,

$$X_i = \begin{cases} m & \text{with probaility } p \\ \frac{1}{m} & \text{with probaility } 1-p. \end{cases} \quad [5]$$

By considering the logarithm of this we get that the process $\log X_t$ is given by,

$$\log X_t = \log x_0 + \sum_{i=1}^{t} \log X_i,$$

that is, $\log X_t$ is an ARW, starting at $\log x_0$, with step size and probability parameters $\log m$ and p respectively. We also construct a corresponding N-step MRW, X_t^N, with time step $\frac{t}{N}$, initial value x_0 and $z_1^N = \frac{1}{z_2^N}, p^N$ as defined below. This formulation preserves the geometric mean and standard deviation of the distribution of X_t (Appendix A).

$$p^N = \begin{cases} 0.5 & \text{if p=0.5} \\ \frac{1}{2}\left(1 \pm \sqrt{\frac{t(2p-1)^2}{4p(N-1)(1-p)+1}}\right) & \text{if p} \neq 0.5 \end{cases} \quad [6]$$

$$z_1^N = \begin{cases} m^{\sqrt{\frac{t}{N}}} & \text{if p=0.5} \\ m^{\frac{t(2p-1)}{N(2p^N-1)}} & \text{if p} \neq 0.5 \end{cases} \quad [7]$$

Again we can take logarithms and convert into an ARW. We are interested in the limiting process, as the time step tends to zero, $Y_t = \lim_{N\to\infty} \log X_t^N = \lim_{N\to\infty} \sum_{i=0}^{N} Y_{i,N}$. Where $Y_{i,N} = log X_i^N$.

In the limit, the distribution of the ARW at time t has mean $(2p-1)t \log m$ and variance $4tp(1-p)\log^2 m$ (Appendix B). The limiting process is the sum of iid RVs with finite expectation and variance, hence, by the Central Limit Theorem,

$$Y_t \sim N\left((2p-1)t\log m, 4tp(1-p)\log^2 m\right).$$

B.1. Rate of convergence to a normal distribution. Hence we should consider for various values of N how close the distribution of Y_t is to a normal distribution, that is, the rate of convergence of the central limit theorem.

Theorem B.1 *Berry-Esseen (11) Let $S_n = \sum_{i=1}^{n} Z_i$ be the sum of n iid RVs with $\mathbb{E}(Z_1) = 0$, $\mathbb{E}\left(Z_1^2\right) = \sigma^2$ and $\mathbb{E}\left(\left|Z_1^3\right|\right) < \infty$ then,*

$$\left|\mathbb{P}\left(\frac{S_n}{\sqrt{n\sigma^2}} \leq x\right) - \Phi(x)\right| \leq \frac{3\mathbb{E}\left(\left|Z_1^3\right|\right)}{\sigma^3\sqrt{n}}$$

Where Φ is the CDF of a standard normal distribution

We can apply this theorem by setting $Z_i = Y_{i,N} - \mathbb{E}(Y_{i,N})$. This allows us to calculate the upper bound.

$$\begin{aligned} \mathbb{E}\left(\left|Z_1^3\right|\right) &= \mathbb{E}\left(\left|\left(Y_{1,N} - \mathbb{E}(Y_{1,N})\right)^3\right|\right) \\ &= \mathbb{E}\Big(\Big|Y_{1,N}^3 - 3\mathbb{E}(Y_{1,N})Y_{1,N}^2 - \\ &\quad \mathbb{E}(Y_{1,N})^2 Y_{1,N} - \mathbb{E}(Y_{1,N})^3\Big|\Big) \\ &\leq \mathbb{E}\left(\left|Y_{1,N}^3\right|\right) + \mathbb{E}\left(\left|3\mathbb{E}(Y_{1,N})Y_{1,N}^2\right|\right) + \\ &\quad \mathbb{E}\left(\left|\mathbb{E}(Y_{1,N})^2 Y_{1,N}\right|\right) + \mathbb{E}\left(\left|\mathbb{E}(Y_{1,N})^3\right|\right) \end{aligned}$$

where we have used the triangle inequality in the final step above. Now we have $\mathbb{E}\left(\left|Y_{1,N}^3\right|\right) = \log^3 m_n$, $\mathbb{E}\left(\left|Y_{1,N}^2\right|\right) = \log^2 m_n$, $\mathbb{E}(|Y_{1,N}|) = \log m_n$, $\mathbb{E}\left(Y_{1,N}^3\right) = (2p_n - 1)\log^3 m_n$.

Further, the components in the equation above become,

$$\begin{aligned} \mathbb{E}\left(\left|3\mathbb{E}(Y_{1,N})Y_{1,N}^2\right|\right) &= \left|3(2p^N - 1)\log m_n\right| \mathbb{E}\left(\left|Y_{1,N}^2\right|\right) \\ \mathbb{E}\left(\left|\mathbb{E}(Y_{1,N})^2 Y_{1,N}\right|\right) &= \left|log^2 m^N\right| \mathbb{E}(|Y_{1,N}|) \\ \mathbb{E}\left(\left|\mathbb{E}(Y_{1,N})^3\right|\right) &= \left|\mathbb{E}(Y_{1,N})^3\right| \end{aligned}$$

Combining this all together gives,

$$\mathbb{E}\left(\left|Z_1^3\right|\right) = \log^3 m^N \left(4\left|2p^N - 1\right| + 2\right)$$

Therefore, given an $N < \infty$, Z_i as defined above and a time t we have that,

$$\left|\mathbb{P}\left(\frac{\sum_{i=1}^{N} Z_i}{\sqrt{N\sigma^2}} \leq x\right) - \Phi(x)\right| \leq \frac{3}{\sqrt{N}} \frac{\left|2p^N - 1\right| + 2}{\left(4p^N(1-p^N)\right)^{\frac{3}{2}}}.$$

This is shown numerically in Figure 3.

The Berry-Esseen Theorem represents a worse-case scenario, Figure 2 shows how closely the CDF of an ARW at time t matches with a normal distribution for different values of n.

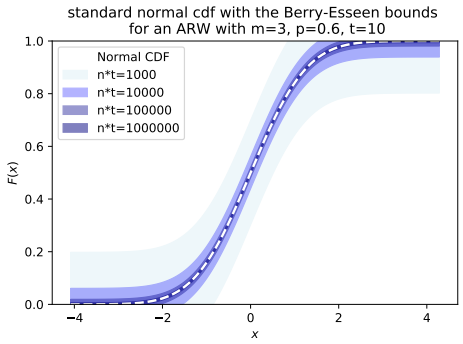

Fig. 1. Berry-Esseen bounds change as n differs

We can see in the final graph of Figure 2 the empirical distribution is outside the Berry-Esseen bounds, however this is due to variation between the empirical and true CDF.

B.2. Limiting ARW is a Brownian Motion. Standard Brownian motion, B_t, is a stochastic process with the properties,

1. $B_t \sim N(0,t)$
2. B_t has stationary, independent increments
3. B_t has continuous paths

To show that our limiting process Y_t is a Brownian Motion we define $B_t = \frac{Y_t - \mu t}{\sigma}$ where $\mu = (2p-1)\log m$ and $\sigma^2 = 4p(1-p)\log^2 m$. Our goal is to show that B_t satisfies the definition above, then $Y_t = \sigma B_t + \mu t$ is a Brownian motion with drift μ and infinitesimal variance σ^2.

1. We have chosen μ and σ in such a way that $B_t \sim N(0,t)$
2. An ARW is a Markov process hence B_t has independent increments, properties of sums of the normal distribution show B_t has stationary increments.
3. Consider $B_t - B_{t_0}$ for $t_0 < t$. Again using properties of sums of the normal distribution, $B_t - B_{t_0} \sim N(0, t - t_0)$. Therefore as $t \to t_0$, $B_t - B_{t_0} \to \delta(0)$ where $\delta(\cdot)$ is the Dirac delta function. Hence $\lim_{t \to t_0} B_t = B_{t_0}$ and thus B_t has continuous paths.

Hence B_t is a standard Brownian Motion and therefore the limit of our ARW is a Brownian Motion.

B.3. Limiting MRW is a GBM. Writing $Y_t = \sigma B_t + \mu t$, the limit of our MRW is therefore given by,

$$X_t = e^{\sigma B_t + \mu t}$$

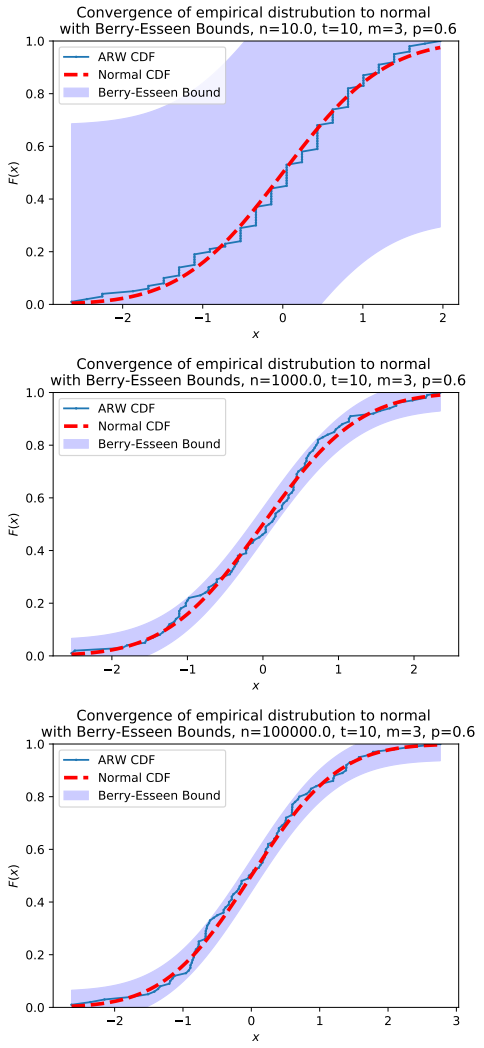

Fig. 2. Showing Berry-Esseen bounds for $n = 10, 1000, 10000$ at $t = 10$, and comparing with the convergence to normal of empirical distribution

Hence the limiting process of an MRW with jump and probability parameters m,p respectively is a GBM with parameters $\mu = (2p-1)\log m$ and $\sigma^2 = 4p(1-p)\log^2 m$.

Note that the by considering the contrapositive we know that the convergence of the ARW implies the convergence of the corresponding MRW. If the process X_t, as defined above, does not converge then $Y_t = \log X_t$ does not converge either. This is because if $X_t \to \infty$ then $\log X_t \to \infty$. Therefore if $Y_t = \log X_t$ converges then X_t also converges.

We have already discussed the differences between the distribution of a MRW and the log-normal distribution (the distribution of GBM). The issues we discussed earlier are overcome in the limit as we tend towards a process with $p = \frac{1}{2}$ and $z_1 = z_2 = 1$ and the maximums agree. The distribution of the MRW converges to log-normal as we expect. However, for any N-step approximation the two maximums and hence the moments, do not agree.

C. Considerations when simulating GBM. This leads to two main considerations when using simulations to investigate systems that include GBM or random multiplicative processes in general.

C.1. Number of simulations required to obtain accurate moments. We have seen that the distribution of a MRW, even when it converges to GBM, does not follow the log-normal distribution. In fact even the first moments do not agree. Further the error in the moments is of exponential order. The reason for this is the dominant role that rare events play in moments. This contrasts with the sample moments which for small samples will be dominated by the most probable events. The underlying principles can be seen in a MRW, to get an accurate estimate of the mean for an N-step MRW you need an order e^N sample size. These issues are compounded further when simulating GBM as any N-step approximating MRW is by construction bounded unlike GBM itself.

C.2. Appropriate Statistics. In light of these issues, it is clear that in many settings using classical moments alone will not give a clear picture of the system. There are alternatives such as the geometric moments that can give more useful information. For example, the geometric mean is precisely the most probable event in a MRW. It is pertinent to consider which is appropriate for the context that you are working within. Particularly in real-world multiplicative regimes, where typical behaviour is often more relevant than the mean.

3. Reallocating geometric Brownian motion

Following (4) (8), to better model the dynamics of wealth throughout the population we now include a social structure to the model through adding a reallocation term. At every time step all agents put a proportion of their wealth into a central fund τx_i and receives back an amount equal to all other agents $\tau \langle x \rangle_N$, where N is the number of agents. When this mechanism is translated into continuous time we get the following model,

$$dx_i = \underbrace{x_i\left(\mu dt + \sigma dW_i\right)}_{GBM} - \underbrace{\tau\left(x_i - \langle x \rangle_N\right) dt}_{reallocation} \qquad [8]$$

The parameter τ controls the proportion of reallocated wealth. Hence when $\tau = 0$ there is no reallocation and the model is a collection of independent GBM processes. In the case $\tau > 0$ the reallocation acts like a tax, and hence becomes a mean-reverting process similar in nature to the Ornstein-Uhlenbeck process (12). When $\tau < 0$ the reallocation term causes agents to repel from the mean, this means agents can have negative wealth even if they start with positive wealth.

We now further investigate the differences and similarities in behaviour between the three different cases for τ.

A. Model behaviour under zero reallocation.

A.1. Scaled distribution which produces stationarity. When $\tau = 0$ the reallocation term vanishes so it is simply a GBM process. Hence at time t the wealth of agents follows a Log-Normal distribution with parameters $(\mu - \frac{\sigma^2}{2})t$ and $\sigma^2 t$. These both depend on t and confirm that the process diverges to ∞. To achieve a stationary distribution we would require neither term to depend on t.

The scaling property of the Log-Normal distribution says that for $X_t \sim Lognormal((\mu - \frac{\sigma^2}{2})t, \sigma^2 t)$ then $X_t e^{at} \sim Lognormal((\mu - \frac{\sigma^2}{2} + a)t, \sigma^2 t)$. Hence even by scaling we still have a t dependent term for the second parameter. Therefore even with scaling we get a process which diverges and so no stationarity.

We are interested in finding out whether the process can be rescaled to a stationary distribution as this is equivalent to the level of inequality becoming constant over time. We can see above that when there is no reallocation in the model the variance between agent's wealth grows over time, however if a stationary distribution exists then the percentage of wealth that the top 1%, 10%, 50% etc. of agents control will stay constant over time. Therefore the fact that no stationary distribution exists for zero reallocation indicates that over time inequality will grow.

A.2. Convergence rates to Limiting Distribution. The result above of the ensemble average $\langle x(t) \rangle_N$ at time t converging to a Log-Normal distribution holds for $N \to \infty$. In this section we will perform a numerical study to decide on an appropriate value of N in our future simulations such that our ensemble means are approximately Log-Normally distributed.

We simulated a Reallocating geometric Brownian motion (RGBM) process with $\tau = 0$ with differing values of N, all agents were given unit initial wealths. Using maximum likelihood estimation we fitted the wealth of agents at time $t = 10$ to a Log-Normal distribution and recorded parameter values. We repeated this to get a distribution for the μ and σ parameters, shown in Figure 3.

We can see from the figure that as N increases the distribution of estimated parameters tends towards a delta function at $(\mu - \frac{\sigma^2}{2})t$ for the μ parameter and $\sigma\sqrt{t}$ for the σ parameter,

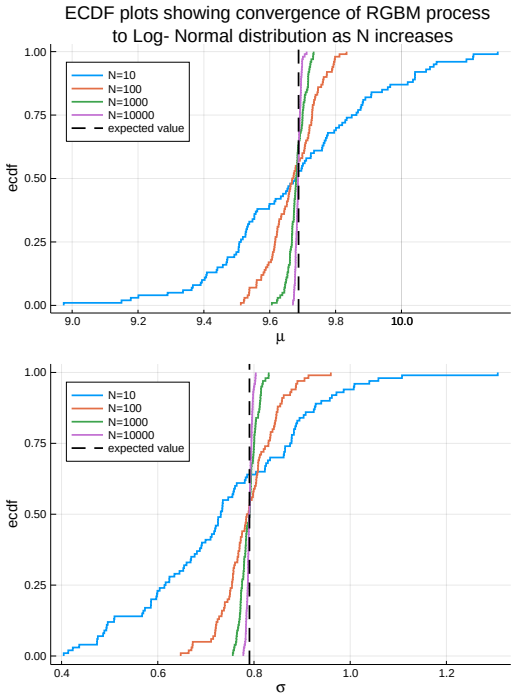

Fig. 3. ECDF plots for the distribution of estimated parameters for a Log-Normal distribution for varying N. The RGBM model was run with parameters $\mu = 1$, $\sigma = 0.25$, $\tau = 0$, $t = 10$ with a sample size of 100.

which matches theoretical results for a GBM process.

For future simulations, in order to balance both computational time and accuracy the value $N = 1000$ was chosen. We will also use this value of N when $\tau \neq 0$ as due to the dependence between agents any extreme behaviour will be removed through reallocation and hence the rate of convergence to results when $N \to \infty$ will be quicker.

A.3. Influence of Initial Distribution (Fixed initial values). We now look at the effect of the initial distribution of agents has on the stationary distribution of the RGBM model. In the case where all agents start with the same wealth we can use the basic scaling properties of the Log-Normal distribution from Section A.1 to show that scaling by $a > 0$ results in the nth moment of the process to increase by a factor of a^n.

For $a < 0$ we can compare the process X_t with initial values $X_0 = x_0$ to the process X' with initial values $X'_0 = -x_0$. At time $t = 0$ we have $\langle x_0 \rangle_N = -\langle x'_0 \rangle_N$ and $x_0 = -x'_0$. Therefore $dx_i = -dx'_i$ at time $t = 0$. Further, if we couple the processes such that they share the same sequence of random variation (through dW_i) then it is clear $x_i = -x'_i$ for all time t. That is, multiplying the initial value by -1 and running the process is identical to running the process and then multiplying all points by -1. In fact this scaling property holds for all $\tau \in \mathbb{R}$.

A.4. Influence of Initial Distribution (Variable initial values). For the case where we have an initial distribution we can write the RGBM model as a mixture-process,

$$Y_t = S_0 X_t$$

where S_0 is our initial distribution and X_t is the GBM process. Moments are given by,

$$\mathbb{E}\left(Y_t^n\right) = \mathbb{E}\left(S_0^n\right) e^{n\mu t + \frac{\sigma^2}{2}(n^2-n)t} \qquad [9]$$

Therefore in order for all moments of the mixture-GBM process to match for two different initial distributions then all their moments must match.

So the initial distribution partly determines the ensemble distribution of the process at time t, but it is not clear whether the ensemble distribution is still Log-Normal. Firstly, if our initial points are spread over positive and negative values then the distribution at time t will no longer follow a Log-normal distribution. This is because the support of the Log-Normal is $(0, \infty)$ and as agents are unable to switch from negative to positive values we are guaranteed to have points outside the support at time t. Further numerical simulation shows that even if the points are shifted by the minimum value so that all points fall inside the range $(0, \infty)$ the distribution is still not Log-Normal. This is most likely because the negative trajectories tend towards $-\infty$, that is, in the opposite direction to the positive trajectories.

Restricting our initial distributions to only take positive values solves the issue of agents' wealth diverging to $-\infty$, however the ensemble distribution of the mixture-process still is not necessarily Log-Normal. Only in the case of our initial distribution S_0 being Log-Normal does the mixture process Y_t retain its Log-Normal ensemble distribution property.

We now consider specifically the case where $S_0 \sim N(\mu_n, \sigma_n^2)$. Fix our time point t and write $X_t = e^{\mu + \sigma Z}$ where $Z \sim N(0, 1)$. Hence,

$$Y_t | Z \sim N\left(\mu_n e^{\mu+\sigma Z}, \sigma_n^2 e^{2(\mu+\sigma Z)}\right)$$

Therefore the distribution of Y_t, $f_{Y_t}(y_t)$, is given by

$$f_{Y_t}(y_t) = \int_{-\infty}^{\infty} f_{Y_t,Z}(y_t, z)dz$$
$$= \int_{-\infty}^{\infty} f_{Y_t|Z}(y_t|z) f_Z(z)dz$$
$$= \int_{-\infty}^{\infty} (2\pi e^{\mu+\sigma z}\sigma_n)^{-\frac{1}{2}} \exp\left(\frac{-\left(y_t - \mu e^{\mu+\sigma z}\right)^2}{2e^{2(\mu+\sigma z)}}\right)\Phi(z)dz$$

Where $\Phi(z)$ is the pdf of a standard normal random variable. This integral has no known analytical form, however as $f_{Y_t|Z}(y_t|z) \leq (2\pi e^{\mu+\sigma z}\sigma_n)^{-\frac{1}{2}}$ by dominated convergence we can conclude this integral is finite.

Using Eq. (9) we can see the moments of Y_t and X_t grow at different rates, hence the distribution of Y_t exists and is not Log-Normal.

A.5. Analysing the process backwards. Analysing the process backwards, that is as $t \to 0$, if at time t the process Y_t is following a Log-Normal distribution then $\mathbb{E}(Y_t) = Ce^{\mu t} \to C$ as $t \to 0$. Further, $\mathbb{V}\text{ar}(Y_t) = De^{2\mu t}(e^{\sigma^2 t} - 1) \to 0$ which shows that in order for the process to reach a Log-Normal distribution, all agents must start with the same wealth. Note that this doesn't contradict our earlier claim that having an initial distribution which is Log-Normal results in convergence to a Log-Normal distribution, this case is the same as recording the process from some time $t' \neq 0$. Using the substitution $t \to t - t'$ in the equations for $\mathbb{E}(Y_t)$ and $\mathbb{V}\text{ar}(Y_t)$ confirms this.

A.6. Summary. To conclude, we have shown that with no reallocation the RGBM model cannot be scaled to produce a stationary distribution, indicating that inequality will grow over time. Scaling the initial condition causes the process to scale at the same rate, but having the initial condition as a distribution results in the process losing the ensemble Log-Normal distribution property, unless that initial distribution is also Log-Normal.

We will now see how these results change when positive reallocation is added into the model.

B. Model behaviour under positive reallocation. When $\tau > 0$, the mean ensemble wealth at time t, denoted $\langle x(t)\rangle_N$, is proportional to $e^{(\mu-\frac{\sigma^2}{2})t}$. However for an individual agent their expected value at time t is proportional to $e^{\mu t}$. In fact, (4) shows that when $t < \frac{\log N}{\sigma^2}$, the per-capita wealth $e^{\mu t}$ is a valid approximation for the rate of growth of $\langle x(t)\rangle_N$. This behaviour is shown in Figure 4.

As a result of the differentiation between the time average and

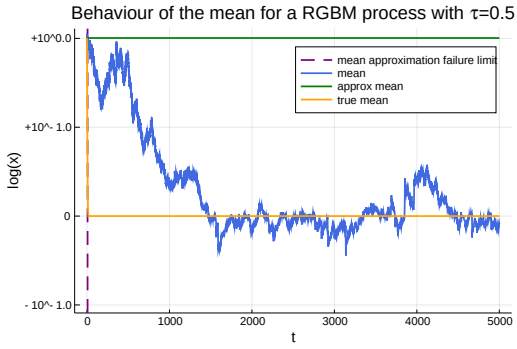

Fig. 4. Mean of RGBM process with parameters $\mu = 0$, $\sigma = 1$, $N = 1000$, $\tau = 0.5$. The value before which the approximation $e^{\mu t}$ for the mean is valid is $t = 6.91$, shown by the vertical purple line.

ensemble average, looking at an agent over time is different to looking at all agents at a single time. Therefore an agent will not necessarily spend 1% of their time in the top 1%, 10% of their time in the top 10% etc. In other words our model has an in-built inequality mechanism, even though all agents start with the same wealth. Formally this is known as non-ergodicity and holds no matter the value of τ in the reallocation term.

B.1. Growth of the mean. The generator of the process, for any τ, is given by,

$$\mathcal{L}f(x_i) = \Big(\mu x_i + (\langle x\rangle_N - x_i)\tau\Big) f'(x_i) + \frac{x_i^2\sigma^2}{2} f''(x_i)$$

Setting $f(x_i) = x_i$ gives,

$$\mathcal{L}f(x_i) = \mu x_i + (\langle x\rangle_N - x_i)\tau$$

Dynkin's Formula (13) on the time evolution of the mean states that $\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}(f(x_i)) = \mathbb{E}(\mathcal{L}f(x_i))$, hence,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}(x_i) = \mu\mathbb{E}(x_i) + \tau\mathbb{E}\left(\langle x\rangle_N - x_i\right). \qquad [10]$$

The first term on the right hand side of Eq. (10) describes the growth of x_i as we would expect in the case with no reallocation. Considering the second term, if $x_i > \langle x\rangle_N$ then the agent's trajectory is above the mean and the value of the second term in Eq. (10) is negative, forcing the trajectory towards the mean. The opposite behaviour occurs when $x_i < \langle x\rangle_N$, that is, the agent's trajectory is above the mean and is forced towards the mean.

This result for the evolution of the mean Eq. (10) also helps to explain the behaviour in Figure 4. Initially all agents have the mean wealth and so the mean grows at rate $e^{\mu t}$. However as

time progresses, as the difference in wealth between agents at the top and agents at the bottom reaches orders of magnitude, the effect of higher wealth agents being pushed towards mean via reallocation has much greater effect on the mean that the effect of lower wealth agents. Thus explaining the reduction in the growth of the mean over time under positive reallocation.

B.2. Rescaling the process to get stationarity. The importance of this is as $N \to \infty$ scaling the process by the per-capita wealth results in the ensemble wealth following an Inverse-Gamma distribution with parameters $\alpha = 1 + \frac{2\tau}{\sigma}$ and $\beta = \frac{2\tau}{\sigma}$ (4). As neither parameter depends on t, the rescaled process is stationary.

Inequality is therefore constant over time, which suggests that the introduction of taxes into a system has the effect of eliminating the growing inequality issue we saw when there was no reallocation.

Additionally the Inverse-Gamma distribution has power-law tails, as does the Log-Normal distribution, hence with the reallocation term added into the model we have kept the same distribution of wealth in the upper tail but stopped it from growing unproportionally large over time.

C. Effect of initial distribution.

C.1. Influence of Initial Distribution (Fixed initial values). If $Y \sim InvGamma(\alpha, \beta)$ then $aY \sim InvGamma(\alpha, a\beta)$. Hence our original process, $X_t \sim InvGamma(1 + \frac{2\tau}{\sigma}, e^{\mu t}\frac{2\tau}{\sigma})$.

Further, using this scaling property, we have that multiplying the initial values by a constant and running the process produces the same result as running the process and then multiplying all points by the constant.

In the case where the scaling constant is negative and the initial condition positive, the distribution of the ensemble wealth is an Inverse-Gamma distribution reflected on the line $x = 0$. As a result agent trajectories tend to $-\infty$ as $t \to \infty$.

When $\tau > 0$ agents can switch from positive to negative wealth and vice versa. Manipulation of Eq. (8) shows that the process with initial values x_0 and drift value $\mu < 0$ when multiplied by -1 is identical to the process with initial values $-x_0$ and drift value μ.

C.2. Influence of Initial Distribution (Variable initial values). We now consider the case where the initial values follow a distribution S_0. The RGBM model becomes a mixture model defined by $Y_t = S_0 X_t$. Considering moments gives,

$$\mathbb{E}\left(Y_t^n\right) = \mathbb{E}\left(S_0^n\right)\mathbb{E}\left(X_t^n\right) \quad [11]$$

$$= \mathbb{E}\left(S_0^n\right)\frac{\frac{2\tau}{\sigma}^n e^{n\mu t}}{(\alpha-1)(\alpha-2)...(\alpha-n)} \quad [12]$$

Which shows that the mixture distribution does not necessarily follow an Inverse-Gamma distribution. However, for large t the component of the moment from the RGBM process dominates the component from the initial distribution, this suggests that the mixture process tends towards an Inverse-Gamma distribution as $t \to \infty$. This proposition was confirmed with numerical simulations.

The question now becomes what parameter values does the distribution of the mixture process converge to? We ran numerical simulations with initial distributions being Normal with different variances, the distribution for the Inverse-Gamma parameter estimates is shown in Figure 5.

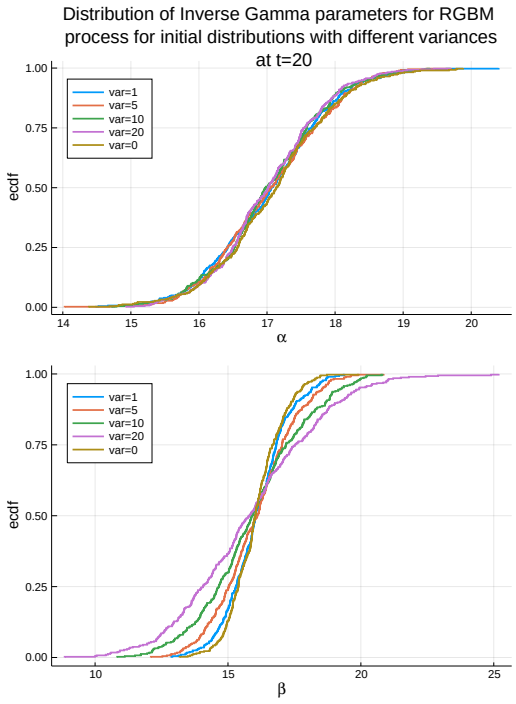

Fig. 5. ECDF plots for the distribution of estimated parameters for an Inverse-Gamma distribution given Normal initial distributions with different variances. The RGBM model was run with parameters $\mu = 0.2$, $\sigma = 0.25$, $\tau = 0.5$, $t = 20$ with a sample size of 400. The limit of validity for the approximation of t is 110.5. The initial distributions all had $\mu = 1$.

We see in Figure 5 that no matter the variation of the initial distribution the distribution of parameter estimates for α by $t = 20$ are almost identical. Figure 6 shows that by time $t = 200$ we can say the same for the distribution of our estimates for β. This confirms our conjecture that as $t \to \infty$ the mixture process tends in distribution to the standard RGBM process with initial value equal to the mean of the

initial distribution.

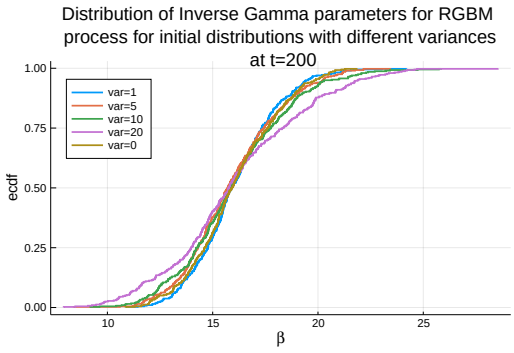

Fig. 6. ECDF plots for the distribution of estimated β parameter for an Inverse-Gamma distribution. The process is identical to that which produced Figure 5 except the distribution is recorded at time $t = 200$ rather than $t = 20$

The rationale for the argument of asymptotic convergence relies on the existence of higher moments in Eq. (11). Analytically a mixture process with a Cauchy initial distribution cannot converge to an Inverse-Gamma distribution as a Cauchy distribution has infinite variance and so $\mathbb{E}\left(Y_t^n\right) = \infty$ for $n > 1$. Numerical simulation of the mixture process with initial distribution $S_0 \sim Cauchy(1, 2)$, shows that the distribution of parameter estimates of β has much heavier tails than when the initial distribution was normally distributed. However performing a Kolmogorov-Smirnov test on the mixture process with a Cauchy initial distribution does not reject the null hypothesis that the ensemble distribution comes from an Inverse-Gamma distribution. This suggests that our choice of $n = 1000$ in this case was not large enough to generate enough extreme values in our initial Cauchy distribution.

C.3. Divergence direction of Process. When $\tau > 0$ all agents tend towards either ∞ or $-\infty$ as $t \to \infty$. Further, due to the mean attraction property of the reallocation term, they all diverge in the same direction. We now investigate the factors that determine the direction that the RGBM process diverges.

For initial distributions with mostly positive initial values the process diverges to ∞. The opposite occurs when the process has mostly negative initial values. Hence we consider the case where the initial distribution is centred around zero.

Using numerical simulation, we considered two initial distributions, Uniform and Normal, both with $\mu = 0$ and $\sigma^2 = \frac{1}{3}$. We recorded the mean of simulated initial distributions and the direction of divergence of the resulting process. Instances were then grouped into bins depending based on their initial mean values, results are shown in Figure 7.

From Figure 7 we can see that the mean of the initial distribution has no significant effect on the divergence direction of

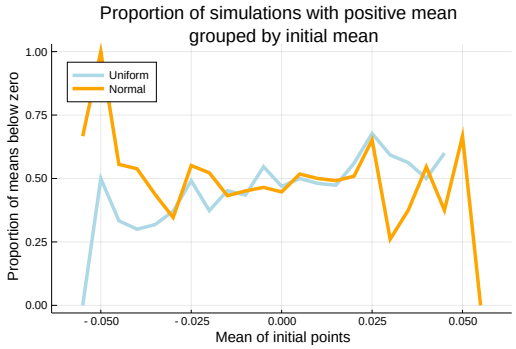

Fig. 7. Plot showing how the mean of the initial distribution affects the state of divergence of a RGBM process with $\tau = 0.5$. Other model parameters were $\sigma = 0.25$, $\mu = 1$, $N = 1000$.

the model, in the case where the mean is close to zero. We tested two other metrics: the proportion of agents starting with negative wealth and the ratio of the variance between agents with positive wealth and agents with negative wealth. Neither produced any significant results. A possible conclusion is that the randomness of the process is the deciding factor for initial distributions around zero.

C.4. Summary. When positive reallocation is added to the model, which in the real-world usually takes the form of taxes, we see the existence of a stationary distribution. This indicates that inequality is kept constant over time, however the non-ergodicity of the process means that although the rich will not get proportionally richer over time they are likely to stay rich for longer than is 'fair'.

Further, when initial values follow a distribution, as long as the distribution has finite moments, the ensemble distribution of that process will converge asymptotically towards distribution of a process with identical initial values for all agents.

When the parameters of the model were estimated using real-life wealth data, Berman et al. (4) found that the direction of reallocation has been negative since the 1980s, this is motivation for us to consider the dynamics of wealth with negative reallocation.

D. Model behaviour under negative τ. For negative reallocation it is easier to use the analogy of agents putting in a portion of their wealth and getting back an equal proportion in the opposite direction. In this scenario agents all put an equal amount of capital into the central fund and receive back an amount proportional to their wealth. Hence, over time, the rich get richer and the poor get poorer.

More specifically, under negative reallocation, the reallocation term causes agents trajectories to be repelled away from the

mean, and hence trajectories split into two strands. In the case $\mu > \tau$, the trajectories tend towards ∞ and $-\infty$ and when $\mu < \tau$ the strands converge together. The behaviour of the process after the point of convergence depends on the parameters, examples of possible behaviours are shown in Figure 8. For the majority of this section we will only consider the process with initial value fixed at 1 for all agents.

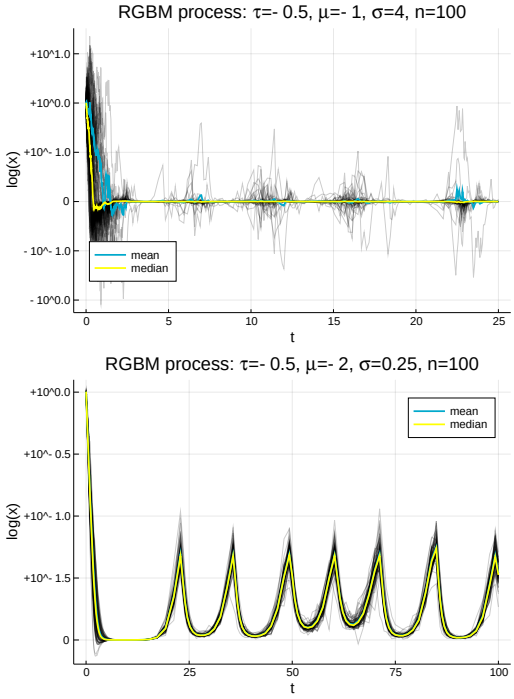

Fig. 8. RGBM processes under the negative τ regime when $\mu < \tau$. Parameters are given in the title of each plot.

In the top plot of Figure 8 we see clusters of volatility. This is caused by a growing feedback process, as $\mu < 0$ agents with a small positive wealth, when the mean is zero, are turned to a slightly larger negative wealth. This process repeats, eventually the trajectories get to a level where the mean value deviates significantly from zero which acts as a balancing mechanism to reduce the volatility back down.

In the bottom plot of Figure 8 we see a periodic half-wave as the long-term behaviour of the process. In this case the effect of μ dominates the effect of σ. In fact as $\sigma \to 0$ we get the same behaviour. It is not clear what causes this. Using Eq. 10, and making the assumption that $\langle x_i \rangle_N = x_i$, which for σ small is approximately correct, we have, $\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}(x_i) = \mu\mathbb{E}(x_i)$ which explains the downwards swoop of the half-wave but not the upwards movement or the sharp turn. We can only conclude from this that at some point one trajectory is far enough from $\langle x \rangle_N$ to have an significant effect on the mean.

We need a high σ to observe this behaviour as variation in the model is what starts the feedback process.

D.1. Growth of the mean. Although strands are repelled from the mean they fall into a state of equilibrium over time where the ensemble mean stays proportionally close to zero and the strands grow at equal rates. This behaviour can be seen from re-writing Eq. (8) as,

$$dx_i = x_i\Big((\mu - \tau)dt + \sigma dW_i\Big) + \tau \langle x \rangle_N dt. \qquad [13]$$

In the long-term, $|x_i(\mu - \tau)| \gg \left|\tau \langle x \rangle_N\right|$, hence by ignoring the final term in Eq. (13) we can approximate the long-term behaviour to the case $\tau = 0$ with μ replaced by $\mu - \tau$. Hence we see that individual trajectories grow at rate $e^{\mu-\tau}$.

The reallocation term causes the rate of growth of the strands to be additive inverses of each other, however due the randomness of the model this balance is not exact and so the mean oscillates. This randomness is proportional to x_i which grows at rate $e^{\mu-\tau}$, hence the ensemble mean also grows at that rate.

The final term in Eq. (13) tells us that swapping between negative and positive wealth can only happen when the sign of a trajectory is the same as sign of mean. These assertions can all be found in Figure 9.

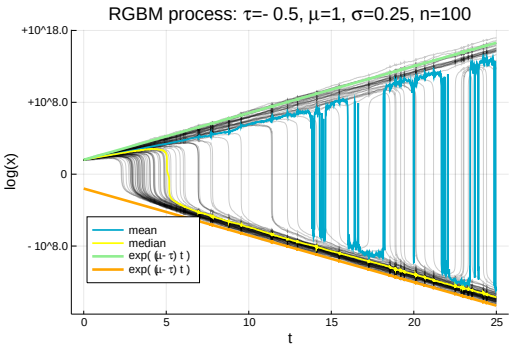

Fig. 9. RGBM process when $\tau < 0$ and $\mu > \tau$. Note that the strands grow at rate $e^{\mu-\tau}$, individual trajectories only cross $y = 0$ in the opposite direction to the ensemble mean at that time and the size and behaviour to the mean relative to the strands (the y-axis is a log scale).

We can also gain information on the behaviour of the mean by considering Eq. (10), which is re-written below,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}(x_i) = \mu\mathbb{E}(x_i) + \tau\mathbb{E}\left(\langle x \rangle_N - x_i\right) \qquad [14]$$

In the negative τ regime the final term on the right hand side of Eq. (14) has the opposite effect as in Section B.1. When $x_i > \langle x \rangle_N$ the trajectory is greater than the mean and hence the final term is positive, so over time (provided the random movements are small) the trajectory moves further away from the mean. Conversely when $x_i < \langle x \rangle_N$, the final term is negative so, again, the trajectory is repelled downwards away

from the mean. This provides another illustration of the mean repulsion mechanism of the reallocation term.

D.2. Linking Long-term behaviour to Geometric Brownian Motion. To validate the assumption that in the long-term we can ignore the final term of Eq. (13), and consider the scenario where we simulate an RGBM process for $\tau < 0$ up to time t_{max}, but at time $t_{div} < t_{max}$ we introduce an RGBM process with $\tau = 0$ and μ replaced with $\mu - \tau$ whose initial conditions are the ensemble distribution of the original process at time t_{div}. We compare the ensemble distributions of both processes at time t_{max}.

Figure 10 shows the results. Applying the Kolomogorov-Smirnov test on both the ensemble distributions of both processes at $t = t_{max}$ fails to reject the null hypothesis that the ensemble distributions come from the same process at the 5% level. We had to choose σ to be small in order to keep the strands approximately independent. Hence for small σ the process is equivalent to when $\tau = 0$ with an initial value given by some distribution and μ replaced with $\mu - \tau$.

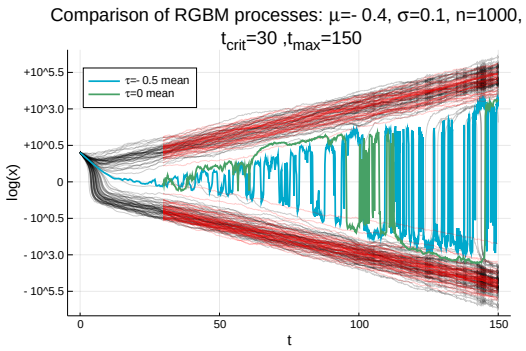

Fig. 10. RGBM process when $\tau < 0$ and $\mu > \tau$. Note that the strands grow at rate $e^{\mu-\tau}$, individual trajectories only cross $y = 0$ in the opposite direction to the ensemble mean at that time and the size and behaviour to the mean relative to the strands.

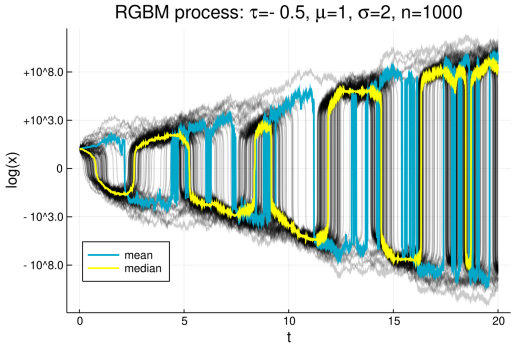

Fig. 11. RGBM process for large σ. The median is part of a group of agents whose trajectories form a wave in time. It is for this reason that we believe in this case a stationary distribution does not exist

D.3. Investigation into Existence of a Stationary Process. For large σ, trajectories clump together and form a wave which suggests as $t \to \infty$ the process never reaches stationarity. This behaviour is shown in Figure 11.

For small σ scaling the process by $e^{\mu-\tau}$ produces a process which numerically is stationary (according to a Kolomogorov-Smirnov test). However, in Section A.1 we noted that when $\tau = 0$ we cannot scale the process to create stationarity. Further in Section D.2 we concluded the process $\tau < 0$ is identical to the process when $\tau = 0$ with an initial distribution which casts doubt on whether the process, when scaled by $e^{\mu-\tau}$, is actually stationary.

Further the higher moments for the rescaled process are plotted in Figure 12, we can see they are not constant over time which is evidence against the idea that scaling the process by $e^{\mu-\tau}$ produces a stationary distribution.

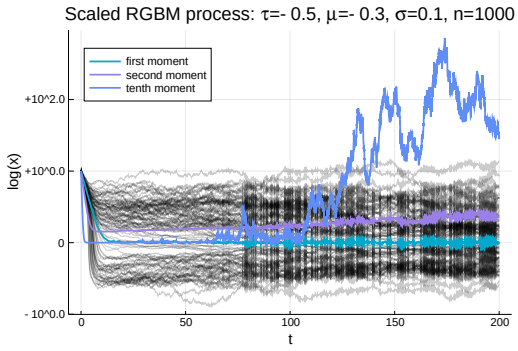

Fig. 12. Scaled RGBM process with various moments. Process parameters are $\mu = -0.4, \tau = -0.5, \sigma = 0.05, t_{max} = 200, n = 1000$.

This puts mathematical evidence behind our negative reallocation mantra that the rich get richer and the poor get poorer, as the non-existence of a stationary distribution means that inequality grows over time.

D.4. Ensemble distribution of process. There is no known result for the ensemble distribution of the RGBM model when $\tau < 0$, but through numerical simulations there are properties we know the distribution must satisfy.

The distribution must have a scaling property, that is, if $X \sim F(\cdot)$ then $aX \sim F(\cdot)$ where $a \in \mathbb{R}$ and F is a distribution. In our case this condition is necessary so that the distribution of the process starting from any initial value will come from the same family of distributions.

Further, the distribution must be symmetric about the y-axis, as in the long-term the rate of growth of negative and positive valued trajectories is equal. This assertion was confirmed numerically.

In an attempt to find the distribution of the RGBM process when $\tau < 0$ we consider the process $y_i = x_i e^{-(\mu-\tau)t}$, due to our required scaling property this is equivalent to finding the distribution of x_i. Using Itô's Lemma on Eq. (8) for y_i with $\langle x \rangle_N = ae^{(\mu-\tau)t}$ for some $a \in \mathbb{R}$, we find that the scaled process follows the SDE,

$$dy_i = \tau a dt + y_i \sigma dW_t.$$

Inputting Eq. (D.4) into the Fokker-Planck equation gives us,

$$\frac{\partial p}{\partial t} = a\tau \frac{\partial p}{\partial y} + \frac{\sigma^2}{2} y\Big(y^2 p\Big).$$

To find the stationary distribution we set $\partial p/\partial t = 0$. Which gives the solution,

$$p(y) = p(0) \exp\left(\int_{-\infty}^{x} \frac{2a\tau - 2\sigma^2 y}{\sigma^2 y^2} dy \right)$$

where $p(0)$ is the boundary condition for $y = 0$. This integral does not converge.

D.5. Number of agents in each strand. The fact that both strands of the process when $\tau < 0$ diverge at the same rate suggests that they contain an equal number of agents. Figure 13 shows the distribution of the number of agents in the upper strand for different values of σ.

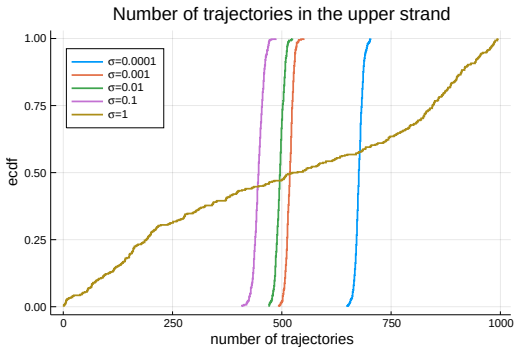

Fig. 13. ECDF plot of the distribution for the number of agents in the upper strand at time $t = 20$ for RGBM process with $\tau = -0.5$, $\mu = 0.2$, $N = 1000$, with a sample size of 400.

For very small σ recording the distribution at $t = 20$ with the chosen parameters is not enough time for the strands to emerge which is why it is higher than for other values. The reason distribution for $\sigma = 0.1$ can be seen in Figure 10 with the wave of trajectories which upsets the 50:50 balance seen for lower values of σ.

D.6. Influence of Initial Conditions on Long-Term distribution. So far we have only considered the process with initial condition $x_0 = 1$ for all agents. As with non-negative τ, multiplying the initial condition by a scalar and running the process is the same as running the process and then multiplying by the scalar.

For initial values that follow a distribution, we have no asymptotic convergence to a mutual distribution. Figure 14 shows that for initial distributions that are normal the distribution at time $t = 20$ is the same as at $t = 60$ and different for each mixture process. This aligns with our results from the $\tau = 0$ process

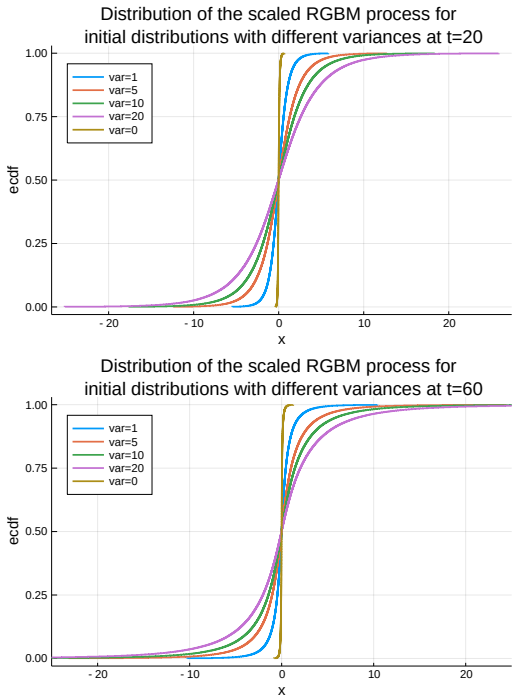

Fig. 14. ECDF plots for the distribution of each mixture process at time $t = 20, 60$ respectively, with Normal initial distributions each with different variances. The parameters of the process are $\mu = -0.2$, $\tau = -0.5$, $\sigma = 0.1$. Using Kolomogorov-Smirnov tests determined that for each initial distribution there is not enough evidence to reject the null hypothesis that the distributions are time $t = 20$,$t = 60$ are the same at the 5% level

D.7. Summary. With negative reallocation, the wealth of agents grows at rate $e^{\mu-\tau}$ either towards ∞ or $-\infty$. In the long-term, and for small values of σ we can approximate the process to one with $\tau = 0$ and an initial distribution. There is no known distribution for the wealth of agents and initial conditions have little effect on the long-term behaviour of the process.

The transition of reallocation from positive to negative has meant the property of constant inequality over time has been lost. This backs the findings of (4) who suggest reallocation in the US has effectively been negative since the 1980s and the report by the World Inequality Lab (6) who show that inequality has been growing since the 1980s.

4. Reallocating GBM with bankruptcy

One way we might be able to improve the model is to add a concept of bankruptcy. In the previous section we proved that in the reallocation model, an individual's wealth may tend to negative infinity. This is not representative of the real world, so we can add a lower bound beyond which an individual's wealth is reset to some value, for example 1 or 0.

For our model, we can take the discrete formulation and add in a step where, at each time step, we check if the value is below the threshold, b, and reset the wealth to 1 it if it is. This model is given by

$$dx_i = x_i(\mu dt + \sigma dW_i) - \tau(x_i - \langle x \rangle_N)dt + \mathbb{I}\{x_i = b\}(1 - b).$$

The indicator function $\mathbb{I}$ is 1 when the event $x_i = b$ is true and 0 otherwise. The model is difficult to examine analytically because sample paths are no longer continuous, so we will examine its behaviour through simulations.

A. Numerical Analysis of the Bankruptcy Model. In the positive τ regime, with positive growth and initial conditions and σ close to what we see in the real world, all trajectories will quickly leave the area where they are impacted by the bankruptcy boundary. If we use negative initial conditions, some trajectories will become more negative until they hit the threshold where they will be reset to 1 and then grow as normal in a positive τ regime. This process is shown in Figure 15. As we are only considering positive growth rates, the negative τ regime will be where we see an impact.

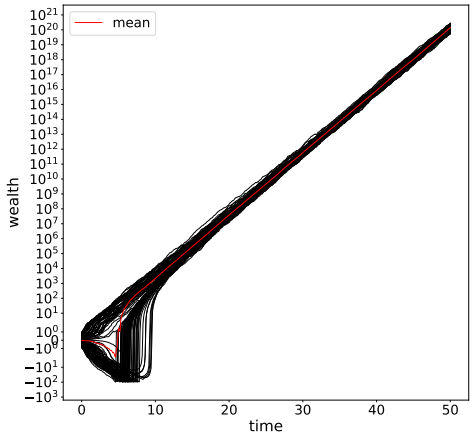

Fig. 15. A graph showing trajectories with initial conditions taken from a uniform distribution between -1 and 1, n = 100, $t_{max} = 10$, $\mu = 1.02$, $\sigma = 0.2$, $\tau = 0.1$, and threshold for bankruptcy at -100. The mean is shown in red.

The behaviour our model shows varies greatly depending on the parameters.

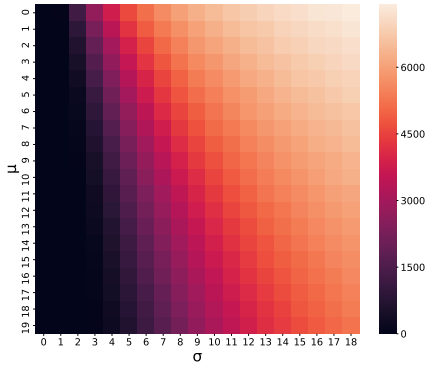

Fig. 16. A figure showing the number of times the threshold at 10^{-5} is hit for varying μ and σ with $\tau = 0.8$. It's clear that only for large σ compared to μ do trajectories reach the boundary.

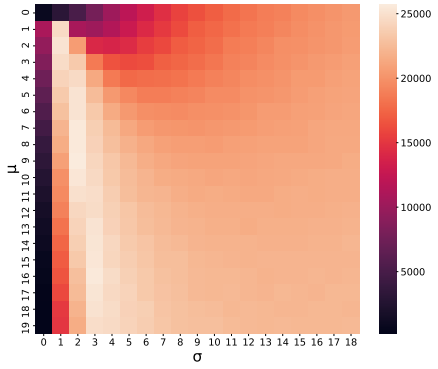

Fig. 17. A figure showing the number of times the threshold at 10^{-5} is hit for varying μ and σ with $\tau = -0.2$. Clearly, it is now very common for the trajectories to reach the boundary.

Figure 17 shows some structure for small μ. We can explore this by simulating trajectories for interesting values of μ and σ. Interesting behaviour can be seen by setting $\mu = 1$ and varying σ.

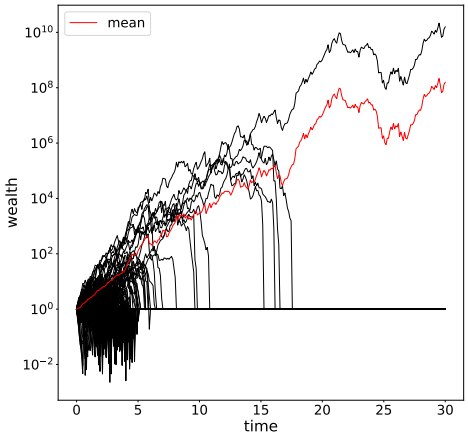

Fig. 18. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 1.0, σ = 1.1 with τ = -0.2. The mean is shown in red.

Figure 18 shows that for small σ some trajectories fall below the mean and are then reduced at each time step by the reallocation term. This is the same mechanism that occurs in the model without bankruptcy. Once this has happened, as the mean is controlled by the largest terms, these trajectories are usually too far below the mean for the noise to be able to help them recover.

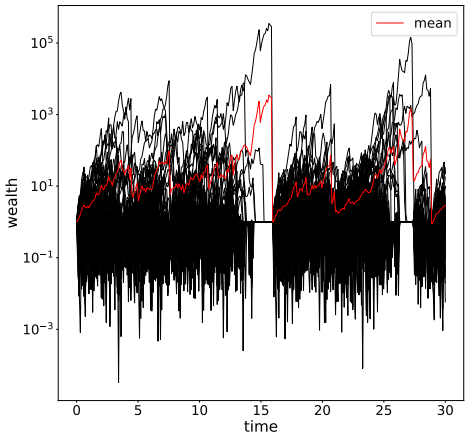

Fig. 19. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 1.0, σ = 2.1 with τ = -0.2. The mean is shown in red.

As shown in Figure 19, for larger σ, no one trajectory can win because the noise is constantly changing where the mean is, and can still bring the highest trajectory below the mean and cause another trajectory to become large.

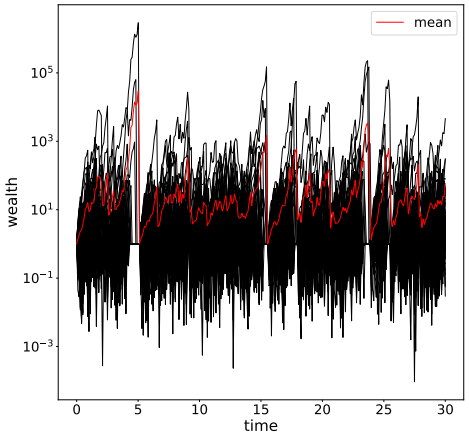

Fig. 20. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 1.0, σ = 4.1 with τ = -0.2. The mean is shown in red.

Figure 20 shows that for very large σ, the trajectories are noise dominated. Trajectories hit the bankruptcy boundary very often as the noise can bring any trajectory below the mean.

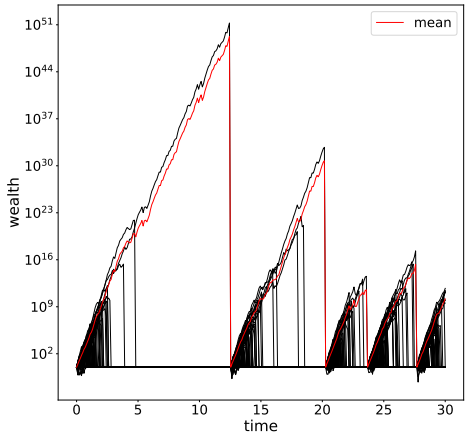

Fig. 21. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 18.0, σ = 6.1 with τ = -0.2, and dt = 0.1. The mean is shown in red.

For the large μ case, as shown in Figure 22 the growth term allows an agent to gain enough wealth so that they are not usually controlled by the noise term, as with small μ and small σ. However, because σ is large, every so often the largest trajectory can be drastically reduced, due to the fact the reallocation term is based on the previous time step. When the noise takes the largest trajectory beneath the mean of the previous time step, it is then brought close to 0 by the reallocation term. Then the process starts again.

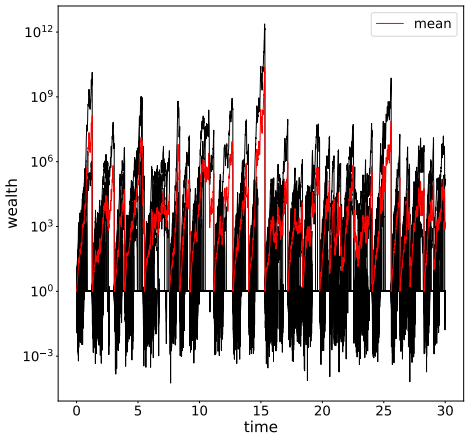

Fig. 22. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 18.0, σ = 6.1 with τ = -0.2, and dt = 0.01. The mean is shown in red.

This simulation used the Euler-Marayama scheme, which only uses values from the previous time step to calculate the next.

Decreasing the size of the time step causes the value of the largest trajectory to fall below the mean more frequently, which is further evidence that the phenomenon is due to the fact that the mean is taken from the previous time step. Therefore this is an artifact of approximating a continuous process numerically.

It might be possible to decrease these effects by choosing a scheme with lower error, such as the Runge-Kutta method.

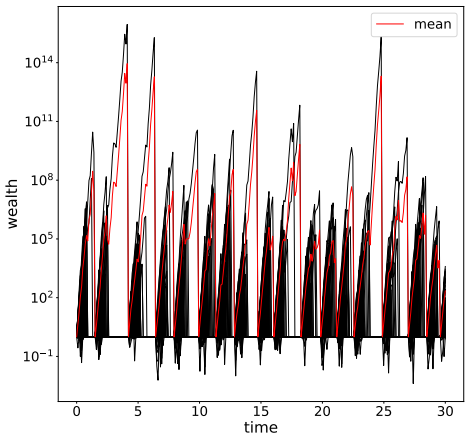

Fig. 23. A graph showing simulated trajectories with n = 100, t_{max} = 30, μ = 18.0, σ = 19.1 with τ = -0.2. The mean is shown in red.

For large μ and large σ, the process is very similar. However, as the noise term is much larger in amplitude, it is able to drag trajectories below the mean much more often, and therefore we have a similar process but much faster than the one shown in Figure 22, as shown in Figure 23.

Now, using values of μ and σ which are more accurate to the real world, ($\mu = 1.02$ and $\sigma = 0.2$ as in Berman et al. (9)), we can find the long term distribution of the system. The time at which this distribution is reached depends on τ. For $\tau = -0.1$ the distribution is usually reached by $t_{max} = 200$, but for $\tau = -0.8$ the distribution is usually reached by $t_{max} = 80$. Of course, the largest value is still increasing in size, therefore there is no stationary distribution. However, the rescaled stationary distribution shows $n - 1$ trajectories at 0, and 1 trajectory at 100.

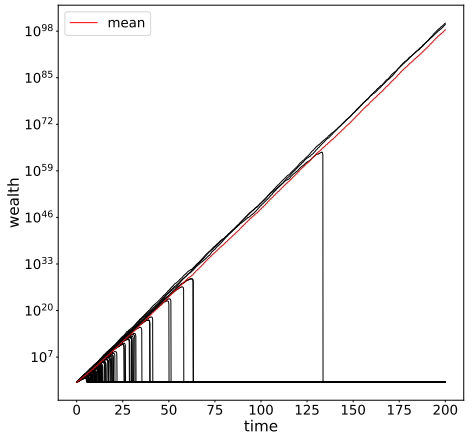

Fig. 24. A graph showing trajectories with n = 100, t_{max} = 200, μ = 1.02, σ = 0.2, and τ = -0.2 (left) with initial conditions between -1 and 1. The mean is shown in red.

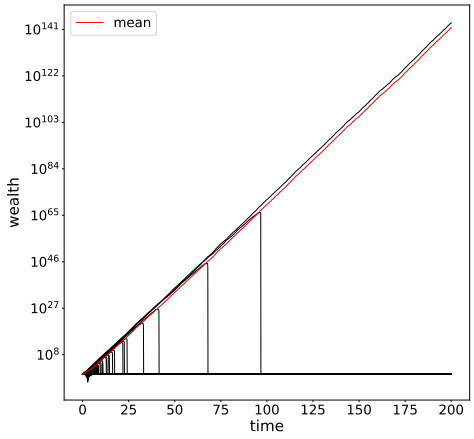

Fig. 25. A graph showing trajectories with n = 100, t_{max} = 200, μ = 1.02, σ = 0.2, and -0.8 (right). The mean is shown in red.

Further analysis shows that the trajectories that appear to collect around 1 in Figure 25 are in fact exactly 1, not just close to 1. This is because, once the long term distribution is reached, the reallocation term will always cause all but one of the trajectories to drop below the threshold for bankruptcy, where they are then all reset to 1. The other trajectory which is large continues to grow.

B. Growth of the Mean. Interestingly, as shown in Figure 26, these simulations do not appear to grow at rate $e^{\mu t}$ or $e^{(\mu-\tau)t}$. In fact, it appears to grow as $e^{(\mu-\tau/2)t}$. However, as shown in 27, changing the time step from 0.1 to 0.01 causes the mean to become closer to $e^{(\mu-\tau)t}$. This is another example of why we must be careful when simulating processes.

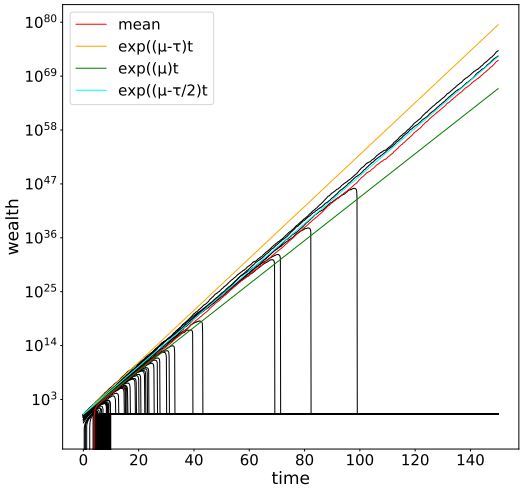

Fig. 26. A graph showing trajectories with n = 100, t_{max} = 100, μ = 1.02, σ = 0.2 and τ = -0.5, and threshold at -100 and dt = 0.1.

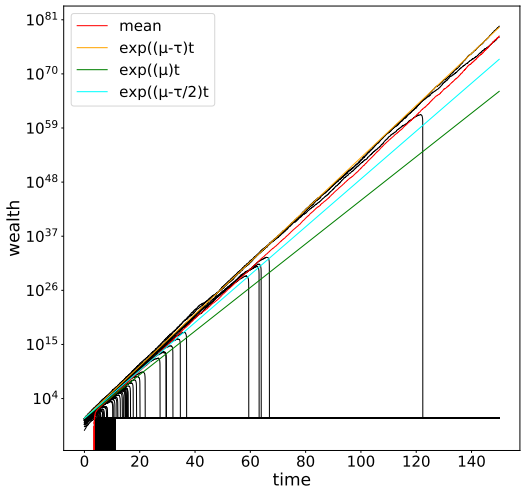

Fig. 27. A graph showing trajectories with n = 100, t_{max} = 100, μ = 1.02, σ = 0.2 and τ = -0.5, and threshold at -100, and dt = 0.01.

C. Varying the Initial Conditions and Bankruptcy Threshold.

Similar behaviour occurs if we use different boundary conditions or lower the threshold for bankruptcy.

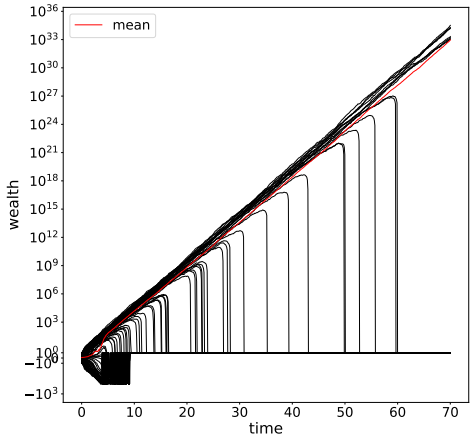

Fig. 28. A graph showing trajectories with n = 100, t_{max} = 70, μ = 1.02, σ = 0.2 and τ = -0.8, with initial conditions drawn from a uniform distribution between -1 and 1, and threshold at -100. The mean is shown in red.

Figure 28 reveals how agents with negative initial conditions quickly approach the bankruptcy threshold. After they are reset to 1 they still have much less wealth than the mean, which by this point is near 10^3. Again these trajectories are pushed downwards by the reallocation term, until the reallocation is bringing them below the threshold at each time step, so that they are always being reset to 1.

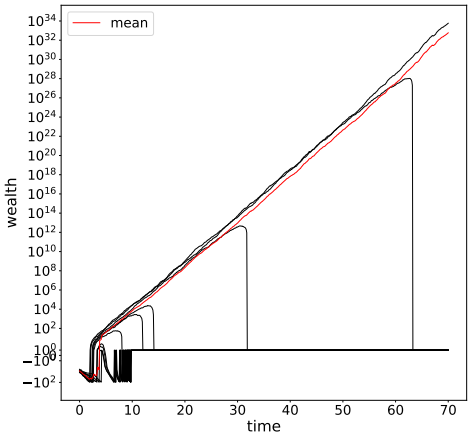

Fig. 29. A graph showing trajectories with n = 10, t_{max} = 70, μ = 1.02, σ = 0.2 and τ = -0.8, with initial conditions drawn from a uniform distribution between -1 and -10, and threshold at -100. The mean is shown in red.

Figure 29 shows that when agents all begin with negative wealth, a few of the largest are able to become positive and compete, while the rest hover between the bankruptcy threshold and the reset value before eventually showing the same behaviour as in previous examples and remaining at 1.

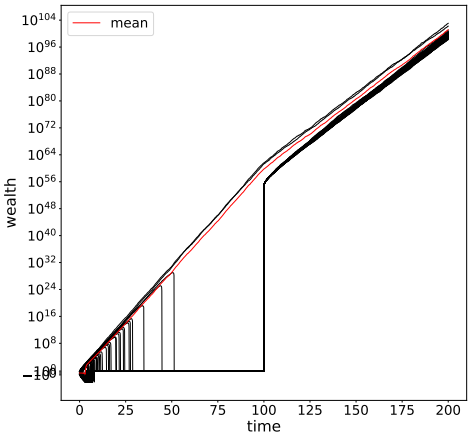

Fig. 30. A graph showing trajectories with n = 100, t_{max} = 200, μ = 1.02, σ = 0.2 and τ = -0.5, between $t = 0$ and $t = 100$, and $\tau = 0.0001$ for the rest of the simulation, with initial conditions drawn from a uniform distribution between -1 and 1, and threshold at -100. The mean is shown in red.

D. Varying τ. Figure 30 shows the behaviour of agents when τ is changed to become positive after the stationary distribution is reached. Even for very small τ, in this case 0.0001, recovery for those who were previously bankrupt is very fast. Although there is a decrease in growth of the mean, the median is clearly much higher than when τ was negative.

E. Evaluation. It is clear that the behaviour shown by this model, in the negative τ regime, is analogous to the behaviour

of the reallocation model without bankruptcy. In this case, rather than one person gaining all of the wealth and the rest of the trajectories going to negative infinity, one person gains all of the wealth and the rest of the trajectories are sustained by the support of bankruptcy. The bankruptcy threshold means that there is no conservation of wealth. Rather than individuals falling into debt, the debt is absorbed by some outside source, for example the government.

In the real world, this may be representative of individuals on a welfare program such as food stamps or Universal Credit, where individuals are sustained but not able to develop savings as they will then lose their support.

One artifact of our model is that the amount a person contributes to the tax pool is proportional to how far they are from the mean, not how much wealth they have. No one should be able to contribute more than they have, and therefore a better model might include a maximum on how much each individual can contribute.

Furthermore, discrete simulations won't have the same long term behaviours as the continuous time model. As no one can contribute more than they have, no one can drop below 0. However, in the discrete time simulations, as the tax is determined by the previous time step, it is possible to fall below 0 and therefore contribute more than you have. At each time step, someone who has been reset to 1 should only be able to contribute $1 \times \tau$ to the tax pool. This may slow down the individual with the largest wealth when many other individuals have low wealth.

In future, we would like to be able to test our first bankruptcy model using data, using similar methods as in Berman et al. (9) in order to see whether or not we have been successful in improving their model. Unfortunately, this is not possible at this time due to a lack of data. We would also like to further explore the behaviours of the second bankruptcy model where we include the idea that an individual cannot contribute more than they have.

5. Summary

We have explored a simple model of wealth dynamics with reallocation between agents. Such models are of growing relevance as wealth inequality increases within society. This included investigating new parameter regimes and introducing a novel concept of bankruptcy via resetting. We also reviewed the connection between a multiplicative random walk and a geometric Brownian motion, and outlined the implications for simulating multiplicative processes. The exploration was primarily through simulation with analytic work when appropriate.

For the model without bankruptcy there are three main regimes. One of the key questions for each regime is whether there is an appropriate re-scaling that produces a stationary distribution. With no reallocation there is no single stationary distribution and the long term distribution is dependent on the initial one. One case to take note of is when every agent starts with equal wealth, the process follows a log-normal distribution. This suggests that in a society where wealth is multiplicative, which it appears to be, an equitable distribution of wealth requires some structure to society. Reallocation can be considered a first attempt to provide such structure, however even with positive reallocation the correctly re-scaled distribution is asymptotically inverse gamma for any initial distribution.

The bulk of the work on the model without bankruptcy was done in the negative reallocation regime, as recent research suggests this is common within the real world. In this case there is no known re-scaling which produces a stationary distribution and empirical evidence suggests one does not exist. We also provide an analytic justification of sign-swapping for the mean and agents, and discussion of novel oscillations induced in the system with low noise when the growth term is less than the reallocation term.

Bankruptcy added a new, more realistic aspect to the model. However, there is a trade off as wealth begins to be generated by bankruptcies and the reallocation term. In addition, it becomes harder to simulate as sample paths are no longer continuous. Given positive reallocation the model behaves very similarly to the model without bankruptcy as very few agents hit the threshold. When reallocation is negative we found that with low-noise eventually only one individual dominates the regime while the other agents stay near or below the reset value. Our investigations suggest that bankruptcy itself does not induce more equity to the system. Significantly, it highlighted the fact that with or without bankruptcy the model can swap between the behaviour associated to different reallocation regimes in relatively short time frames when the parameter is changed. This may be relevant when considering policy changes that would affect reallocation in the real world.

This investigation provides a better understanding and extension of RGBM, which highlights key and novel areas for further analytic and empirical research.

ACKNOWLEDGMENTS. The authors would like to acknowledge the contributions of Prof Colm Connaughton, as well as their funding body the ESPRC.

1. M Marsili, S Maslov, Y Zhang, Dynamical optimization theory of a diversified portfolio. *Phys. A* **253**, 403–418 (1998).
2. B Francis-Devine, Poverty in the uk: Statistics. *House Commons Libr.* (2020).
3. Magazine, Monitor, Who, what, why: What is the gini coefficient? *BBC, Available at https://www.bbc.co.uk/news/blogs-magazine-monitor-31847943* (2015).
4. Y Berman, O Peters, A Adamou, Wealth inequality and the ergodic hypothesis: Evidence from the united states. *Available at SSRN 2794830* (2019).
5. BK Chakrabarti, A Chatterjee, S Yarlagadda, *Econophysics of Wealth Distributions*. (Springer-Verlag Italia), (2005).
6. F Alvaredo, L Chancel, T Piketty, E Saez, G Zucman, The world inequality report 2018 (https://wir2018.wid.world/) (2017).
7. UNDP, Human development report 2014: Sustaining human progress - reducing vulnerabilities and building resilience (http://hdr.undp.org/en/content/human-development-report-2014) (2014).
8. J Bouchaud, M Mezard, Wealth condensation in a simple model of economy. *Phys. A* **282**, 536–545 (2000).
9. Y Berman, O Peters, A Adamou, Far from equilibrium: Wealth reallocation in the united states. *Available at https://arxiv.org/abs/1605.05631v1* (2016).
10. S Redner, Random multiplicative processes: An elementary tutorial. *Am. J. Phys.* **58** (2002).
11. W Feller, *An introduction to probability theory and its applications.* (Wiley), (1971).
12. GE Uhlenbeck, LS Ornstein, On the theory of the brownian motion. *Phys. review* **36**, 823 (1930).
13. B Oksendal, *Stochastic differential equations: an introduction with applications.* (Springer Science & Business Media), (2013).

6. Appendix

A. Scaling of MRW step size and probabilities. Given a set of values $A_1, ..., A_N$ the geometric mean and standard deviation is defined as,

$$\mu_g = \left(\prod_{i=1}^{N} A_i\right)^{\frac{1}{N}}$$

$$\sigma_g = \exp\left(\sqrt{\frac{\sum_{i=1}^{N}\left(log(\frac{A_i}{\mu_g})\right)^2}{N}}\right)$$

For an MRW as specified in Eq. (5), with initial value $X_0 = 1$, the geometric mean and standard deviation over a single time step is,

$$\mu_g = m^{2p-1} \quad [15]$$

$$\sigma_g = \exp\left(\sqrt{p\log^2\left(\frac{m}{\mu_g}\right) + (1-p)\log^2\left(\frac{1}{\mu_g m}\right)}\right) \quad [16]$$

$$= \exp\left(2\sqrt{1-p}\sqrt{p}\log m\right) \quad [17]$$

Considering the same process over N steps, the position of the process after these N steps is $m^i \times \frac{1}{m}^{N-i} = m^{2i-N}$ where i is the number of upwards movements and hence follows a binomial distribution. Let $p_i = \mathbb{P}(Z = i)$ where $Z \sim Bin(N, p)$, then we can write the geometric mean of our N step MRW as,

$$\mu_g = \prod_{i=0}^{N}(m^{2i-N})^{p_i} = (m^{2p-1})^N \quad [18]$$

A similar calculation gives the n-step geometric standard deviation,

$$\sigma_g = \exp\left(\sqrt{\sum_{i=1}^{N} p_i\left(log(\frac{m^{2i-N}}{\mu_g})\right)^2}\right) \quad [19]$$

$$= \exp\left(2\sqrt{1-p}\sqrt{Np}\log m\right) \quad [20]$$

To distinguish between the parameters p, m for the 1-step MRW and the n-step MRW for the remainder of this section we will use the substitution p, m in Eq. (18), Eq. (19) with p^N, z_1^N respectively.

The expression in Eq. (6) can we found by equating the expressions in Eq. (15) with their counterparts in Eq. (18), Eq. (19). We split this into the two cases,

1. $p = \frac{1}{2}$. $2p - 1 = 0 \implies m^{2p-1} = 1 \implies (z_1^N)^{N(2p^N-1)} = 1 \implies N(2p^N - 1) = 0 \implies p^N = \frac{1}{2}$.
   Equating the expressions for the standard deviations gives, $\frac{N}{4}\log^2 z_1^N = \frac{1}{4}\log^2 m \implies \log^2(z_1^N)^{\sqrt{N}} = log^2 m \implies z_1^N = m^{\frac{1}{\sqrt{N}}}$
2. $p \neq \frac{1}{2}$. $(z_1^N)^{2p^N-1})^N = m^{2p-1} \implies z_1^N = m^{\frac{2p-1}{N(2p^N-1)}}$
   Equating the expressions for the standard deviations gives, $(1-p^N)Np^N\log^2 z_1^N = (1-p)p\log^2 m \implies \frac{(1-p^N)p^N}{(2p^N-1)^2} = N\frac{(1-p)p}{(2p-1)^2}$
   Let $A = \frac{N(1-p)p}{(2p-1)^2}$ then the final equation in the above paragraph is equivalent to the quadratic equation, $(4A+1)(p^N)^2 - (4A+1)p^N + A = 0 \implies p^N = \frac{1}{2} \pm \frac{1}{2\sqrt{4A+1}}$
   As $N \to \infty \implies A \to \infty \implies p^N \to \frac{1}{2}$. Similarly $z_1^N \to 1$.
   Hence the $\pm$ sign is $+$ if $p > \frac{1}{2}$ and $-$ if $p < \frac{1}{2}$.

B. Mean and Variance of Y_t. The following contains the calculations for i) $\mathbb{E}\left(Y_{i,N}\right)$, ii) $\mathbb{E}\left(Y_{i,N}^2\right)$, iii) $\mathbb{E}\left(Y_t\right)$, iv) $\mathbb{E}\left(Y_t^2\right)$

1. $\mathbb{E}\left(Y_{i,N}\right)$
$$\mathbb{E}\left(Y_{i,N}\right) = \log z_1^N \times p^N + -\log z_1^N \times (1-p^N)$$
$$= (2p^N - 1)\log z_1^N$$
2. $\mathbb{E}\left(Y_{i,N}^2\right)$
$$\mathbb{E}\left(Y_{i,N}^2\right) = \log^2 z_1^N \times p^N + \log^2 z_1^N \times (1-p^N)$$
$$= \log^2 z_1^N$$
3. $\mathbb{E}\left(Y_t\right)$
$$\mathbb{E}\left(Y_t\right) = \lim_{N\to\infty}\mathbb{E}\left(\sum_{i=1}^{N} Y_{i,N}\right)$$
$$= \lim_{N\to\infty}\sum_{i=1}^{N}\mathbb{E}\left(Y_{i,N}\right)$$
$$= \lim_{N\to\infty} N \times (2p^N - 1)\log z_1^N$$
When $p = 0.5$, $N \times (2p^N - 1)\log z_1^N = 0$ otherwise, assume w.l.o.g $p > 0.5$,
$$N \times (2p^N - 1)\log z_1^N = N \times (2p^N - 1) \times \frac{t(2p-1)}{N(2p^N-1)}\log m$$
$$= (2p-1)t\log m$$
Hence $\mathbb{E}\left(Y_t\right) = (2p-1)t\log m$, this expression matches the expectation when $p = 0.5$

4. $\mathbb{E}\left(Y_t^2\right)$

$$\begin{aligned}
\mathbb{E}\left(Y_t^2\right) &= \lim_{N\to\infty} \mathbb{E}\left(\left(\sum_{i=1}^{N} Y_{i,N}\right)^2\right) \\
&= \lim_{N\to\infty} \sum_{i=1}^{N}\sum_{j=1}^{N} \mathbb{E}\left(Y_{i,N}Y_{j,N}\right) \\
&= \lim_{N\to\infty} N\mathbb{E}\left(Y_{1,N}^2\right) + N(N-1)\mathbb{E}\left(Y_{1,N}Y_{2,N}\right) \\
&= \lim_{N\to\infty} N\mathbb{E}\left(Y_{1,N}^2\right) + \lim_{N\to\infty} N(N-1)\mathbb{E}\left(Y_{1,N}\right)^2
\end{aligned}$$

Now, $\lim_{N\to\infty} N\mathbb{E}\left(Y_{1,N}^2\right) = \lim_{N\to\infty} N\log^2 z_1^N = 4tp(1-p)\log^2 m$ and

$$\begin{aligned}
N(N-1)\mathbb{E}\left(Y_{1,N}\right)^2 &= N\log^2 z_1^N (N-1)(2p^N-1)^2 \\
&= t4p(1-p)\log^2 m(N-1)(2p^N-1)^2 \\
&\to t^2\log^2 m(2p-1)^2 = \mathbb{E}\left(Y_t\right)^2
\end{aligned}$$

Therefore, $\mathbb{E}\left(Y_t^2\right) = 4tp(1-p)\log^2 m + \mathbb{E}\left(Y_t\right)^2$. Which implies $\mathbb{V}\mathrm{ar}\left(Y_t\right) = 4tp(1-p)\log^2 m$