


An agent-based model of the observed distribution of wealth in the United States

Hunter A. Vallejos¹ · James J. Nutaro²  · Kalyan S. Perumalla²

Received: 1 November 2016 / Accepted: 25 July 2017 / Published online: 3 August 2017
© The Author(s) 2017. This article is an open access publication

Abstract Pareto cautiously asserted that the wealth and income distributions which bear his name are universal, basing his argument on observations of this distribution in many different types of economies. In this paper, we present an agent based model (and a scalable approximation of it) in a closely related spirit. The central feature of this model is that wealth enables an individual to secure more wealth. Specifically, the important and novel feature of this model is its ability to simultaneously produce both the Pareto distribution observed in empirical data for the top 10% of the population and the exponential distribution observed for the lower 90%. We show that the model produces these distributions of wealth when initialized with an equitable distribution.

This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the US Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>). Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the US Department of Energy. Source code for the model is available at http://web.ornl.gov/~nutarojj/wealth_model.zip.

✉ James J. Nutaro
nutarojj@ornl.gov

Hunter A. Vallejos
hunterandresvallejos@gmail.com

Kalyan S. Perumalla
perumallaks@ornl.gov

¹ Georgia Institute of Technology, Atlanta, GA, USA

² Oak Ridge National Laboratory, Oak Ridge, TN, USA

Then, using historical data, we initialize the model with US wealth shares in 1988 and show that the model tracks wealth share changes from 1988 to 2012. Simulations to 2088 project that the top 0.01% of the population will possess more than 70% of the total wealth in the economy.

Keywords Pareto · Boltzmann–Gibbs · Wealth · Income · Inequality · Agent-based model

1 Introduction

The Pareto distribution was first discovered by Vilfredo Pareto in the late 19th century (Chipman 1976; Persky 1992). He found that, regardless of how an economy is structured, the wealth and income distributions among the upper tier of the population (typically, the upper 10% of a population is examined) follow a power law, which came to be known as a Pareto distribution (Jones 2015). The Pareto distribution is given by the relation

$$P(\omega^* > \omega) = \left(\frac{\omega_m}{\omega}\right)^\alpha, \quad (1)$$

where ω_m is some minimum wealth and α is the Pareto index. Equation 1 states that the probability of someone's wealth ω^* being greater than ω is proportional to $\omega^{-\alpha}$. Thus, it is a power law, and, in most economies, α is between 1 and 3 (Ghosh et al. 2016; Gabaix 2009; Chipman 1976).

Rapid improvements in the availability of data on wealth have reinvigorated empirical and theoretical investigation of the Pareto distribution. One recent discovery is that the lower 90% of the population does not follow a power law, but rather follows an exponential distribution¹ (Drăgulescu and Yakovenko 2001), which is defined by the probability density function

$$P(\omega) = ce^{-\omega/T_\omega}, \quad (2)$$

where T_ω is the “wealth temperature” and c is a normalizing constant. These exponential and Pareto distributions fit to observed wealth distributions in the United Kingdom (UK) are illustrated in Fig. 1, which is taken from Drăgulescu and Yakovenko (2001) and Yakovenko and Rosser (2009).

A common measurement of inequality in the distribution of wealth is the Gini coefficient. If we arrange the individuals in an economy by their wealth in m distinct groups such that $\{\omega_i\}_{i=1}^m$ is non-decreasing ($\omega_{i+1} \geq \omega_i$ for all i) and $\Omega = \sum_i^m \omega_i$ is the total wealth of the economy, then the Gini coefficient, G , is defined as:

$$G = \frac{2}{m\Omega} \left(\sum_{i=1}^m i\omega_i \right) - \frac{m+1}{m}. \quad (3)$$

¹ Although the gamma and log-normal distributions have been used in other models, we restrict our study to the exponential distribution, which is the only distribution we have found that matches empirical measurements of the bottom of the wealth distribution.




Fig. 1 Cumulative wealth distribution of UK copied directly from Drăgulescu and Yakovenko (2001)

There have been many attempts to describe why wealth inequality arises (Yakovenko and Rosser 2009). The macro-economic perspective has shifted to Piketty's findings published in his *Capital in the Twenty-First Century* (Piketty 2013). The model that emerges from Piketty's work generates a Pareto distribution through the combined effects of several factors, some of most significant being birthrate, growth rate, and return rate (Jones 2015).

While Piketty's work focuses almost exclusively on the Pareto distribution of wealth, a variety of econophysical models, which rely extensively on descriptions of stock market transactions (Yakovenko and Rosser 2009), either produce a distribution characteristic of the exponential distribution, or a power law, but not both. Moreover, by relying on specific economic data and processes to generate wealth inequality, these models are intrinsically tailored to a specific type of economy. Consequently, these models can explain inequality within a specific economic system, but they do not shed light on Pareto's assertion that inequality is intrinsic to all economic systems.

In contrast, we propose a model that fits the entire wealth distribution without modeling any specific economic process other than growth. The fundamental assumption in this model is that individuals who possess greater wealth derive from it a greater ability to secure more wealth. We quantify this advantage and show that the resulting dynamic in a growing economy leads to a Pareto distribution of the wealthy minority and exponential distribution of the less wealthy majority. The model's free parameter is calibrated to reproduce the observed distribution of wealth in the United States (US) from the period between 1988 and 2012. Projections with this model to 2088 produce a striking separation between the shrinking minority of super-rich and everyone else.

2 Model

We assume an exponentially growing economy

$$\Omega(t) = \Omega(0)e^{\lambda t} \quad (4)$$

where $\Omega(t)$ is the sum of the wealth of all individuals at time t and λ is the growth rate of the economy. In our model, the economy does not grow continuously, but instead grows in small, fixed increments $\delta\omega$. Hence, over a time T , the number of increments n will be

$$n = \frac{\Omega(0)(e^{\lambda T} - 1)}{\delta\omega}. \quad (5)$$

Each increment is given to exactly one individual. Let $\omega_{i,k}$ be the wealth possessed by the i th individual following the k th distribution, with $k = 0$ indicating the initial wealth of i . The probability $\Psi(i, k + 1)$ that i will receive increment $k + 1$ is

$$\Psi(i, k + 1) = \frac{p(\omega_{i,k})}{\sum_i p(\omega_{i,k})}, \quad (6)$$

where $p(\omega_{i,k})$ is a “wealth power.” In our model, we let

$$p(\omega_{i,k}) = (\omega_{i,k})^\beta. \quad (7)$$

Values for λ , $\Omega(0)$, and the initial $\omega_{i,0}$ can be determined directly from data. We assume β to be a parameter intrinsic to the economy, and with it we can directly control the change in inequality. This is quickly explainable by Eq. 7. For $0 < \beta < 1$, the power of wealthier individuals is diminished much more than the power of poorer individuals. For $\beta = 1$, the model gives all individuals a proportional amount of power, and for $\beta > 1$, the model gives a disproportionate amount of power to the richer individuals in the market. We will choose β by fitting economic trajectories produced by the model to historical data.

It is unclear how to select $\delta\omega$, which is primarily a computational artifact and has no obvious analog in the real world. If $\delta\omega$ is on the order of $\omega_{i,0}$, then the first individual to win an increment accrues a tremendous advantage. This circumstance is unlikely to produce a realistic outcome and suggests that $\delta\omega$ should be small. If $\delta\omega$ is much smaller than $\omega_{i,0}$, then those who win wealth increments gain a negligible advantage. Presumably, a very small $\delta\omega$ would eventually approach the same distribution as a slightly larger $\delta\omega$, but our simulations show that the distribution does not correctly form within a reasonable amount of simulated time, and in addition, the computational time implied by Eq. 5 puts a practical lower limit on $\delta\omega$. We have therefore chosen a value of $\delta\omega$ that is small in proportion to $\omega_{i,0}$, but large enough that the individual who receives a wealth packet gains a marginal advantage. A more comprehensive study of how $\delta\omega$ affects outcomes will be reserved for future work.

3 Simulation

We will develop two closely related simulators for the proposed model: one which repeatedly distributes a single wealth packet and updates $\Psi(\cdot, k + 1)$ as described in Sect. 2, and one that distributes multiple packets between updates to $\Psi(\cdot, k + 1)$. The

single-packet model will be used to reproduce the trajectory of wealth shares in the US. The multi-packet model is shown to produce trajectories similar to the single-packet model, but we leaves for future work an investigation of the larger populations enabled by the multi-packet model.

3.1 Single-packet model

The procedure for the single-packet model can be summarized as follows:

1. Set k to 0 and distribute $\Omega(0)$ to the population according to a desired initial distribution to obtain $\omega_{i,0}$ for each individual i .
2. Calculate $\Psi(i, k + 1)$ for each individual within the population.
3. Select at random one individual, say j , in accordance with the distribution defined by the $\Psi(\cdot, k + 1)$.
4. For the selected j , set $\omega_{j,k+1} \leftarrow \omega_{j,k} + \delta\omega$. For all other individuals $i \neq j$, $\omega_{i,k+1} \leftarrow \omega_{i,k}$.
5. Increment k and then repeat from Step 2.

This single-packet version of the model allows for immediate updating of $\Psi(\cdot, k + 1)$ after distributing each wealth packet, which more convincingly suggests that the resulting distribution is an immediate consequence of Eq. 6.

3.2 Multi-packet model

Ideally, we would like to simulate our model with as many individuals as the population of the US. However, Eq. 5 restricts the single-packet model to a smaller population, because the computational time increases nearly quadratically with the number of individuals, m . Thus, we investigate a second version of the model by allowing multi-packet distributions per iteration in order to linearly decrease the computation time. This allows us to simulate larger populations.

Note that, by performing multi-packet distributions, we restrict individuals from immediately using gained wealth to their advantage. Rather they must wait for the next multi-packet distribution. Because of this, it is reasonable to suppose that the dynamics of a multi-packet model would potentially be different than the single-packet model. Nonetheless, for small $\delta\omega$ and large populations, the multi-packet model produces results very similar to the single-packet model in Sect. 4.3.

We employ Algorithm 1 to model the distribution of economic growth via $P \geq 1$ wealth packets per distribution. Note that n in Eq. 5 must be proportionally adjusted to account for the multi-packet distributions, and that the power function in Eq. 7 is only updated after distributing multiple packets.

The algorithm works as follows. In lines 2–3, the wealth of each individual is set to suit any desired initial distribution. After initialization, in line 4, the initial wealth values are sorted in ascending order. The main iteration loop is started at line 6. In each iteration of the main loop, the following four operations are performed.

Algorithm 1: Parallel algorithm for wealth distribution

m	Number of individuals
$W[m]$	$W[i]$ is current wealth of individual i
T	Total time of simulation
1 P	Number of wealth packets per iteration
dw	Value of each wealth packet
R_n	Number of “normal” threads, $1 \leq R_n \leq m$
R_p	Number of “packet” threads, $1 \leq R_p \leq P$

```

2 for  $i = 0$  to  $m - 1$  do
3    $W[i] \leftarrow$  Initial wealth of individual  $i$ 
4 Sort  $W$ 
5  $t \leftarrow 0$ 
6 while  $t < T$  do
7   with  $R_n$  threads in parallel do
8      $\hat{W} \leftarrow W^\beta$ 
9      $S\hat{W} \leftarrow \sum_{i=0}^{m-1} \hat{W}[i]$ 
10     $P\hat{W} \leftarrow$  Prefix sums of  $\hat{W}$ 
11    (prefix sum  $P\hat{W}[i] \equiv \sum_{j=0}^i \hat{W}[j]$ )
12  for  $R_p$  threads in parallel do
13     $g \leftarrow$  Uniform random number in  $[0, 1]$ 
14     $v \leftarrow g \cdot S\hat{W}$ 
15    Bisection search  $P\hat{W}$  for  $j$  such that
16       $j = 0$  if  $P\hat{W}[0] \leq v$ 
17       $0 < j < m$  if  $P\hat{W}[j - 1] \leq v < P\hat{W}[j]$ 
18    Atomic:  $W[j] \leftarrow W[j] + dw$ 
19  Sort  $W$ 
20   $t \leftarrow t + P$ 

```

- (a) *Preparation of the cumulative density function* This is performed in lines 7–11. From the wealth array W , a new array \hat{W} is created to store W^β (line 8). The sum of all elements of \hat{W} is obtained as $S\hat{W}$ (line 9). A new array $P\hat{W}$ is built as prefix sums of \hat{W} . This $P\hat{W}$ is essentially a scaled representation of the cumulative density function of \hat{W} . Locating the position of a random value v in the range $[0, S\hat{W}]$ from the prefix sums array $P\hat{W}$ is equivalent to correctly sampling the cumulative density function. Note that all the steps in this operation can be performed in parallel, affording a parallelism potential of up to $R_n \leq m$ threads.
- (b) *Parallel selection of wealth recipients and their enrichment* This is performed in lines 12–18. Note that each step in this operation will involve action on at least one wealth packet, and hence the parallelism potential of this operation is limited to $R_p \leq P$ threads. Each thread selects a random value v in $[0, S\hat{W}]$ (lines 13–14), and performs a bisection search with time complexity of $O(\log m)$ (lines 15–17) to find the individual who should gain the next packet. The individual’s wealth is updated to reflect this gain – in line 18, the wealth is atomically updated (since theoretically it is possible for more than one thread to happen to choose the same individual within an iteration).

- (c) *Update of the wealth distribution to sorted order* This is performed at line 19.
- (d) Finally, *advancement of time* is performed in line 20 by an increment equal to P packets.

3.3 Simulation scenarios

Our simulations examine two cases. In the first case that we consider, the initial wealth is equally distributed throughout the population. Beginning from this equitable situation, the expected exponential and Pareto distributions form after two centuries. In the second case, the initial distribution of wealth mirrors that in the US in 1988 and we show that, with an appropriate selection of β , the model generates the distribution of wealth observed in 2012.

To fit the model to historical data on the distribution of wealth in the US, we construct exponential and Pareto distributions for the lower and upper parts of the initial distribution respectively. The two distributions are blended using a weighted average over an interval of overlap, where the overlap is defined to span the (normalized) population in (κ_l, κ_u) , with $0 \leq \kappa_l \leq \kappa_u \leq 1$. The resulting distribution of wealth over the whole population is $f_\omega(p)$, which calculates the wealth for a certain percentile p of the population as

$$f_\omega(p) = -T_\omega[1 - \pi(p)] \log\left(\frac{p}{c}\right) + \pi(p)\omega_m p^{-1/\alpha}, \tag{8}$$

where c , T_ω , ω_m and α retain their definitions from Eqs. 1 and 2 and

$$\pi(p) = \begin{cases} 0 & p < \kappa_l \\ \frac{p - \kappa_l}{\kappa_u - \kappa_l} & \kappa_l \leq p \leq \kappa_u \\ 1 & p > \kappa_u. \end{cases} \tag{9}$$

We set $\omega_m = f_\omega(\kappa_l - \epsilon)$ where $\epsilon > 0$ is very small so that the $f_\omega(p) > \omega_m$ for all p in the Pareto part of the distribution.

We use data collected by [Saez and Zucman \(2014\)](#) on the US wealth distribution from 1913 to 2012. From their data, it is apparent that the distribution of the top percentiles has been steadily increasing since 1988. We therefore choose to approximate the interval from 1988 to 2012. This avoids the difficulty of changing β with time, which is likely to be necessary to account for the acceleration in wealth concentration that began in the 1980s.

We initialized the model to match each percentile’s fractional share of the total wealth in 1988. We found that the data in 1988 fit well with a Pareto index of 1.55. The initial conditions for the model are compared with the real US wealth distribution in 1988 in Table 1. We then run the model to see if it will track the US wealth distribution over the next 24 years, letting the market growth rate $\lambda = 0.02588$ since that is the average growth in the Gross Domestic Product (GDP) of the US over the 24 year time period ([Bureau of Economic Analysis, U.S. Department of Commerce 2016](#)).

Table 1 Initial wealth share comparison of model to data (Saez and Zucman 2014)

Percentile	Wealth (%) in 1988	Model
Bottom 90%	34.7	34.7
Top 10%	65.3	65.3
Top 5%	50.9	50.6
Top 1%	27.9	27.7
Top 0.5%	21.5	21.2
Top 0.1%	11.6	11.1
Top 0.01%	4.4	3.8




Fig. 2 Market distribution. Parameters: 10,000 individuals over 225 years with $\delta\omega = 0.01$, $\omega_i = 1$, $\lambda = 0.03$ and $\beta = 1.36$

4 Model results

4.1 Single-packet uniform start

Figure 2 shows the distribution of wealth generated by the single-packet model with $\beta = 1.36$, $\lambda = 0.03$, and the initial distribution of $\Omega(0)$ being uniform throughout the population. The log-log plot of the cumulative distribution is shown in Fig. 3, and it is similar in its form to Fig. 1.

Figure 4 compares the wealth of the bottom 90% of the simulated population to an exponential distribution that is fit to the simulation data. Figure 5 shows a Pareto distribution fit to the top 10% of the simulated population. The growing deviation at wealth 2500 in the simulation data in Fig. 4 is due to the simulation data transitioning from the exponential distribution to the Pareto distribution as seen in Fig. 5, similar to the log-log plot in Fig. 1.




Fig. 3 Log-log of cumulative wealth distribution




Fig. 4 Exponential lower part of distribution. Fit: $y = ae^{-x/b}$; $a = 1.583$, $b = 512.777$

Beginning from its equitable initial state, this simulated economy needs about 200 years to form a distribution of wealth comparable to what is typical of economically developed countries today. Its Pareto index is $\alpha = 1.685$ compared to the UK's Pareto index of 1.9 in 1996 (Drăgulescu and Yakovenko 2001).




Fig. 5 Pareto tail of distribution. Fit: $y = mx + b$: $m = -1.762$, $b = 4.614$, $R^2 = 0.997$

4.2 Single-packet historical fit

As shown in Fig. 6, the single-packet model is able to closely track the historical distribution of wealth to 2012 when initialized with wealth shares in 1988. The Gini coefficient in our model is 0.796 in 2000, which is near the 0.801 Gini coefficient of the US in 2000 (Davies et al. 2006). The variability of the metrics in Fig. 6 is less than 0.1% over 25 runs.

Figure 7 shows a continuation of the simulation to 2088, where an accelerating concentration of wealth is apparent. By 2024, the top 0.01% of the population will take possession of as much wealth as the bottom 90% of the population. Moreover, the model predicts that by 2088 the top 0.01% of the population will possess more than 70% of the total wealth, which is the same fraction of total wealth owned by the top 10% in 2012. This implies that the Pareto part of the wealth distribution is shrinking, and it will continue to shrink beyond 2088. Other indicators of this trend are the Gini coefficient and Pareto index that are both approaching 1.

4.2.1 Historical fit from an equitable start

The same simulation described above, when begun from an equitable distribution in which each individual starts with 1 unit of wealth, nearly reproduces the historical data after approximately 360 years for the top 10% and bottom 90%. These data are shown in Fig. 8. It is interesting and unexpected that the distribution of inequality passes through the historical data in spite of the equitable start. However, simulated and historical data for the top 5%, 1%, etc. are not in close agreement. The difference




Fig. 6 Model tracking real wealth distribution data from 1988 to 2012. This is with initial parameters: 100,000 individuals and setting $\beta = 1.25$, $\delta\omega = 0.01$, $c = 1$, $T_\omega = 0.3597$, $\kappa_l = 0.75$, $\kappa_u = 0.9$, and $\alpha = 1.55$. Dotted lines represent the model's fit




Fig. 7 Model Projection from 1988 to 2088




Fig. 8 Model projection begun from an equitable start

between historical wealth ownership starting at 1988 and simulated wealth ownership starting at 360 years for each category is shown in Fig. 9. The magnitude of the differences between the historical data and simulated data grow with time, which indicates that the rate of concentration in the model differs from that observed in reality even if the general trends are in agreement.

4.3 Validation of multi-packet model

To validate our multi-packet model, we ran it with the same parameters as the single-packet distribution model in Fig. 2. Figure 10 shows the power law fit to the top 10% of the population produced by the multi-packet model, with 1024 wealth packets per distribution over a population of 10,000 individuals. The power law created by the multi-packet model is very comparable to the result of our single-packet model shown in Fig. 5, and it produced a very similar Pareto index $\alpha = 1.636$.

Preliminary runs of the model with large numbers of individuals produce the expected distributions as well. We believe that because the dynamical behavior of the multi-packet model is essentially the same as the single-packet model, the multi-packet model will likely produce the same results in Sect. 4.2, and thus would make the same projection about wealth inequality in the US. Large scale simulation experiments are a topic for future work.

5 Discussion of dynamics

The value of β is central to the dynamics of our model, as experimentation with it on the presupposed exponential and Pareto distributions in Sect. 4.2 has shown that it

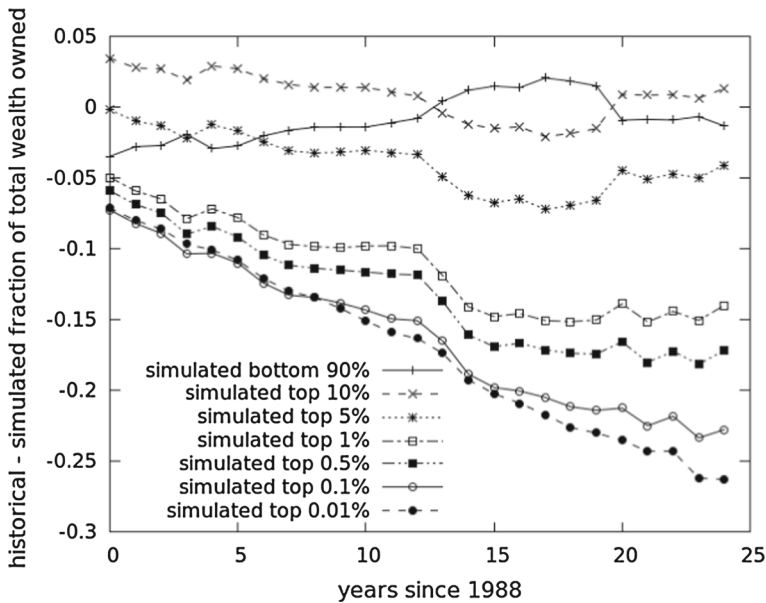


Fig. 9 Difference between model and historical data from an equitable start




Fig. 10 Pareto tail of wealth distribution formed by multi-packet model with 1024 wealth packets per distribution. Each packet is worth $\delta\omega = 0.01$. Fit: $y = mx + b$; $m = -1.636$, $b = 18.849$, $R^2 = 0.994$

controls the change in the wealth inequality in an economy. This is concurrent with the predicted role of β in Sect. 2. For $0 < \beta < 1$, we observed that the value of the Pareto index α would increase, implying less inequality in the population. For

$\beta = 1$, the model would not alter the initial exponential and Pareto distributions for an extended period of time, maintaining its initial value of α . Lastly, with $1 < \beta < 2$, we observed the distribution moving to lower values of α , suggesting more inequality. For any $\beta \geq 2$, we found that the model would quickly diverge from the expected distributions, resulting in a “winner takes all” situation, where there was only one individual far wealthier than the rest of the population. For $1 < \beta < 2$, we observed similar behavior, but the divergence was much slower.

Simulations have shown that $\beta = 1$ is not sufficient to create an exponential and Pareto distributions with $1 < \alpha < 3$ when the model is initialized with a uniform distribution. In most cases, $\alpha > 3$ for $\beta = 1$ when run for a long period of time. Lastly, when projecting from historical data, we found that the part of the population described by the Pareto distribution decreased in size as $\alpha \rightarrow 1$ over time. We attribute this to the divergence of the model, and the closer α was to 1, the more quickly a single individual became extremely wealthy, which is not in accordance with the exponential and Pareto distributions (for $\beta > 1$).

6 Conclusion

The model we propose mirrors empirical wealth distributions in the US. The underlying assumption of the model is in accordance with Pareto’s view that human nature, and in particular our innate tendency to use our resources to improve our own position, is sufficient to generate an uneven distribution of wealth. Indeed, we may hypothesize that economic growth and the tendency for wealth to breed more wealth are central to simultaneously forming the exponential and Pareto distributions of wealth inequality.

The proposed model could also be used to generate income inequality under other appropriate assumptions. In particular, if we assume an increase in GDP is reflected as an increase in income for individuals across the economy, and if the wealth power function is reinterpreted as an income power function, then the model would distribute increases in income as GDP grows. Since there is no change in the dynamics, the model would mimic the exponential and Pareto distributions of income inequality as well.

The historical data considered here span a short period for two reasons. First, there is a scarcity of data quantifying the growth of inequality at other times. Second, from the data that was available to us, the period from 1988 to 2012 was the best example of sustained economic growth. Because sustained growth is a central assumption of the model, it was natural to focus our first validation study on that period. More extensive validation with an extended model that can account for greater economic volatility will be the subject of future research.

A more substantial limitation of the model when trying to reproduce long historical trajectories is its inability to account for social, economic, and geopolitical factors that induce very sudden returns to more equitable conditions. Prime examples of such upheaval in the United States are the Great Depression and the end of the Second World War. It is notable that there are no modern examples of wealth concentrating as our model anticipates for 2088, and this is in spite of the short 76 years that the projection covers (from 2012–2088). This fact seems to limit the model to relatively

short spans between events that alter the dynamics of the economy in some significant way.

Nonetheless, there is a need to understand the parameter β in more depth, as it seems to dictate the rate at which inequality grows or diminishes. One weakness in our approach to selecting β was a focus on fitting the Pareto part of the empirical data. The exponential growth of the economy combined with the preferential attachment of new wealth to already wealthy individuals is expected to produce an essentially exponential distribution of wealth; that is, the exponential part of the distribution is a given. Indeed, the Pareto portion of the simulated distributions shrink with time. Hence, our approach is centered on getting a good fit to the most transient part of the data. An interesting direction for future work would be to reverse this procedure by fitting β to the exponential majority of the data and then seeing how accurately the Pareto minority is matched by the model.

Another direction for future work is to explore what specific economic or human tendencies underlay β . Of the many possibilities, at least two are intuitively appealing. One of these is that the rate at which an individual can consume wealth to maintain a standard of living is almost certainly limited (a point made by Hanauer (2014) when he stated “I earn about 1000 times the median American annually, but I don’t buy thousands of times more stuff. My family purchased three cars over the past few years, not 3000.”). At the lower end of the wealth spectrum, we might expect the vast majority of an individual’s wealth to be expended on basic needs. Hence, these individuals may have less opportunity - less excess wealth - with which to improve their position. At the upper end, the rate limit is reached and excess wealth accumulates, leaving vast reserves to seize new opportunities. Another closely related effect is the redistribution of wealth that occurs through economic activity (e.g., the purchase of a good or service has the effect of transferring wealth from buyer to seller), taxation, charities, and other mechanisms.

Exploring these types of possibilities requires a more refined model of the population to account for the statistical distribution of consumption behavior, engagement in wealth transferring activities, and other factors that vary, possibly substantially, from individual to individual. The scalable algorithm presented in Sect. 3.2 will be an important tool for these future investigations.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bureau of Economic Analysis, US Department of Commerce (2016) NIPA tables (table 1.1.2) section 1 (1969-present) contributions to percent change in real gross domestic product. code: 10102. <http://www.bea.gov/national/nipaweb/DownSS2.asp>, <http://www.bea.gov/national/nipaweb/DownSS2.asp>
- Chipman JS (1976) The paretian heritage. *Revue européenne des sciences sociales* 14(37):65–171. <http://www.jstor.org/stable/40369090>, reprinted with corrections at <http://users.econ.umn.edu/~jchipman/HERITAGE.pdf>

- Davies JB, Sandström S, Shorrocks AB, Wolff EN (2006) The level and distribution of global household wealth. In: NBER Working Paper Series, National Bureau of Economic Research
- Drăgulescu A, Yakovenko VM (2001) Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. *Physica A* 299(1–2):213–221
- Gabaix X (2009) Power laws in economics and finance. *Annu Rev Econ* 1:255–294
- Ghosh A, Chatterjee A, Inoue J, Chakrabarti BK (2016) Inequality measures in kinetic exchange models of wealth distributions. *Physica A* 451:465–474
- Hanauer N (2014) The pitchforks are coming...for us plutocrats. *Politico Magazine*. <http://www.politico.com/magazine/story/2014/06/the-pitchforks-are-coming-for-us-plutocrats-108014?o=0>
- Jones CI (2015) Pareto and Piketty: the macroeconomics of top income and wealth inequality. *J Econ Perspect* 29(1):29–46
- Persky J (1992) Retrospectives Pareto's law. *J Econ Perspect* 6(2):181–192
- Piketty T (2013) *Capital in the twenty-first century*. Harvard University Press, Cambridge
- Saez E, Zucman G (2014) Wealth inequality in the United States since 1913: evidence from capitalized income tax data. In: NBER Working Paper Series, National Bureau of Economic Research
- Yakovenko VM, Rosser JB (2009) Colloquium: statistical mechanics of money, wealth, and income. *Rev Mod Phys* 81:1703. doi:10.1103/RevModPhys.81.1703