Journal of Economic Behavior and Organization 229 (2025) 106820

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Agent-based models of the United States wealth distribution with Ensemble Kalman Filter

Yannick Oswald [a,*], Keiran Suchak [b], Nick Malleson [b]

[a] Institute of Geography and Sustainability, University of Lausanne, Switzerland
[b] School of Geography, University of Leeds, United Kingdom

ARTICLE INFO

Dataset link: https://github.com/yannickoswald/real_time_ineq_abm

Keywords:
Agent-based modelling
Data assimilation
Ensemble Kalman filter
Wealth inequality
Real-time prediction

ABSTRACT

The distribution of wealth is central to economic, social, and environmental dynamics. The release of high-frequency distributional data and the rapid pace of the complex global economy makes 'real-time' predictions about the distribution of wealth and income increasingly relevant. For instance, during the COVID-19 pandemic in spring 2020, the stock markets experienced a crash followed by a surge within a brief period, evidently reshaping the wealth distribution in the US. Yet economic data, when first released, can be uncertain and need to be readjusted — again specifically so during crisis moments like the pandemic when information about household consumption and business returns is patchy and drastically different from "business-as-usual". Our motivation here is to develop one way of overcoming the problem of uncertain 'real-time' data and enable economic simulation methods, such as agent-based models, to accurately predict in 'real-time' when combined with newly released data. Therefore, we tested two distinct, parsimonious agent-based models of wealth distribution, calibrated with US data from 1990 to 2022, in conjunction with data assimilation. Data assimilation is essentially applied control theory — a set of algorithms aiming to improve model predictions by integrating 'real-time' observational data into a simulation. The algorithm we employed is the Ensemble Kalman Filter (EnKF), which performs well in the context of computationally expensive problems. Our findings reveal that while the base models already align well historically, the EnKF enables a superior fit to the data.

1. Introduction

The distribution of wealth is considerably more unequal than that of income or consumption across the population (Chancel et al., 2022). In the United States wealth inequalities are particularly dramatic when compared to other high-income countries. The top 1% of wealth owners in the population own ∼35%, the bottom 50% own approximately nothing at all or only very little wealth. In a typical European country, say France, the top 1% 'only' own 25% of the wealth and the bottom 50% of the population roughly 5% (World Inequality Database, 2024). The circumstances in the USA are partly shaped by a high propensity for individuals to participate in the asset markets and a high degree of dynamism in business and the labour market. More than 60% of all adult Americans own stocks (Jones, 2023) compared to, for instance, only around 18% in Germany (Aktieninstitut, 2023) where wealth inequality is largely shaped through inheritance and low social mobility. The large inequality and dynamism in the US is significant because wealth is an enabler for consumption over time and evidently shapes power-dynamics in business and politics. It also determines environmental outcomes such as CO_2 emissions (Chancel, 2022).

* Corresponding author.
E-mail address: y-oswald@web.de (Y. Oswald).

https://doi.org/10.1016/j.jebo.2024.106820
Received 31 January 2024; Received in revised form 8 November 2024; Accepted 12 November 2024
Available online 22 November 2024
0167-2681/

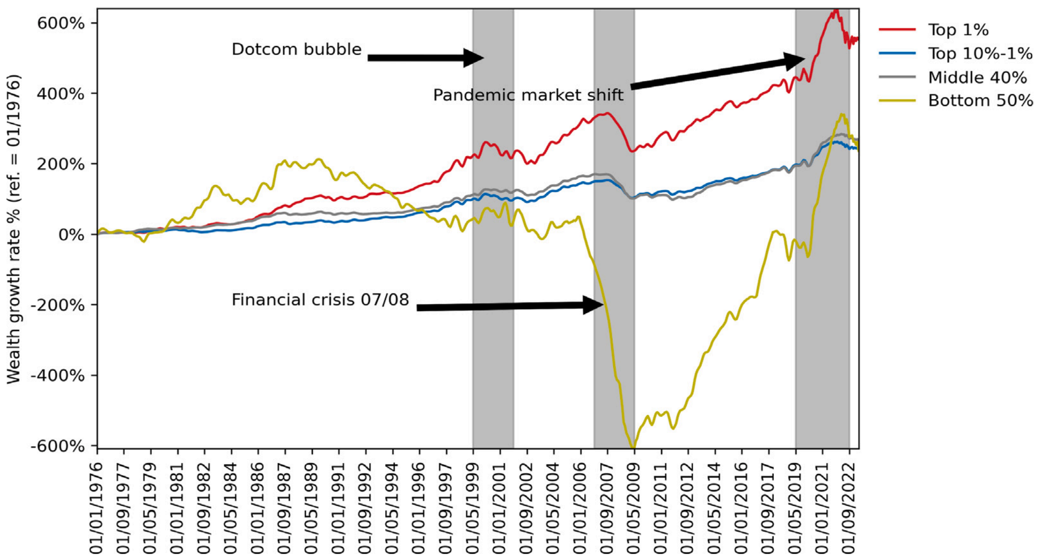

Fig. 1. Wealth distribution in the United States from 1976 to 2022. The data are from https://realtimeinequality.org/ with details in Blanchet et al. (2022).

The distribution of wealth is shaped by political agendas, economic trends and, increasingly, crises. The most recent example for this is the market crash during the early COVID-19 pandemic and the subsequent stock-rally on technology and renewable energy companies. During the 2008 financial crisis, the impact on the wealth distribution in the USA was even more prominent (see Fig. 1). While all groups lost wealth during the financial crisis, the poorest half of the population, measured by wealth, exhibited distinct dynamics and was pushed far below their wealth level of 1976, which for many meant accumulating debt. A lot of poorer and middle-class people defaulted on their mortgages and the stock market crashed too. It took a decade for the Bottom 50% to recover from this. The resilience of the wealth distribution during crisis moments requires improvement and therefore detailed monitoring also on fine-grained time-scales.

Historically, the measurement of wealth, and estimation of its distribution, are cumbersome and slow tasks. Piketty assembled the data for his famous work 'Capital in the 21st century' over many years through historical tax records. Typically, the most detailed wealth data have been released or made available on an annual basis. Recently, however, motivated by the fast-paced crisis dynamics of the pandemic, Blanchet et al. (2022) constructed monthly distributional wealth data — a step towards 'real-time' observation of the wealth distribution. The authors did so based on individual tax accounts and through rescaling of monthly macroeconomic aggregates, such as GDP. Due to the increasing volatility and unpredictability of global economic dynamics, it is of interest to economists and policy makers to predict the distribution of wealth at high frequency in order to counter-steer shock and negative trends. Moreover, yearly and quarterly GDP data are often quite uncertain when first released and regularly need to be readjusted after months, years, or even decades. On top of that, the definitions of macroeconomic metrics change over time (Semieniuk, 2024). And, as with GDP, high-frequency wealth data can be expected to be uncertain and to be readjusted over time. Therefore, observations alone are perhaps not reliable enough on their own to make accurate assessments about the state of the economy, especially so during crisis moments.

At the same time, several computational approaches to economics emerged whose intention is to capture the dynamic complexity of the economy better than traditional equilibrium-focused models do. One such approach is agent-based modelling (ABM). Agent-based models (ABMs) have been identified as one option to go beyond classical analytical economic models in the sense that they are a more natural analog to the discrete and interactive nature of the economy (Axtell and Farmer, 2022) and they have been successfully applied across economic themes from housing markets to financial markets and many more. Historically, a lot of ABM work has remained at the conceptual level. Agent-based models are primarily used to describe stylized facts and overall patterns about isolated phenomena such as how a particular distribution of wealth emerges, how business cycles form and so forth. However, with increasing data availability, recently there is a trend to calibrate ABMs to data in detail and also to employ agent-based models for economic forecasting — one example being on credit risk and inequality by Papadopoulos (2019). With respect to COVID-19, infectious-disease-economy models have been employed to predict pandemic impacts on the economy (Silva et al., 2020). Moreover, there is a new trend towards large-scale macroeconomic agent-based models that aim to replicate entire economies, to construct a digital twin of an economy so to speak. Poledna et al. (2023) constructed such a model for Austria for example, with 8.8 million agents mapping the entire Austrian population and a large number of parameters determining the agent behaviour. Poledna et al. (2023) achieve parity with other standard economic forecasting methods such as state-of-the-art dynamic stochastic equilibrium models (DSGE) and econometric methods in the short-term. However, even such sophisticated models produce forecasts with substantial uncertainty over time and Poledna et al. limit their forecasts to a few years. Data-assimilation in combination with such macro-economic ABMS could perhaps estimate and adjust agent microstates and thus increase the reliability of the predictions further and enable longer prediction horizons. Moreover, although the use of sophisticated ABMs might improve forecasting reliability further, no model will ever be accurate to a degree that is sufficient when facing entirely unforeseen scenarios that are not endogenous to the model, such as the COVID-19 pandemic and its repercussions.

And therefore, due to the inherent uncertainties in both economic data and models, in this work we test the combination of economic agent-based models with data assimilation to overcome these limitations. Data assimilation is a set of algorithms that feed data in real-time into models and aim to optimally combine new observations with model predictions. In other words, it is a Bayesian inference technique that combines a "prior" estimate of a system state (a probability distribution), the model, with new information given, the observations (another probability distribution), to arrive at a posterior estimate. It has a considerable history in meteorology and has contributed to the vast improvements in weather predictions over the last decades (Kalnay, 2003), with some historic applications in engineering and indeed also in economics. In contrast to calibration, which sets *time-fixed* parameters of a model before a forecasting application, data assimilation updates the internal state of a model based on newly available data during the forecast. In terms of agent-based modelling this even can be the *time-varying* agent-specific variables, in our case this will be the wealth per agent. We will apply the Ensemble Kalman Filter (EnKF), which is a variant of the Kalman Filter — a popular tool from optimal control — designed to specifically handle ensembles of model simulations. Most economic agent-based models are stochastic in character and produce model ensembles for forecasting purposes, so this type of data assimilation seems intuitively appropriate.

We pursue two goals in this paper:

(i) we study two distinct, parsimonious agent-based models, one borrowed from the literature, and one developed from scratch, and calibrate them to the USA wealth distribution over time;
(ii) for the first time, we integrate dynamic agent-based models of economic inequality with data assimilation as a proof-of-concept and in particular test the performance of the Ensemble Kalman Filter (EnKF).

Here we do not make use of an elaborate, high-dimensional model such as the one described in Poledna et al. (2023) for several reasons. First, as we will see, the American wealth distribution can be reproduced by very simple, low-dimensional agent-based models with a respectable degree of accuracy. It would have been possible to choose a more accurate and complicated model to simulate the wealth distribution, and indeed some of the assumptions in our case study models are somewhat unrealistic. Secondly, parsimonious agent-based models still play an important role in economic research and elucidate core principles and mechanisms of complex economies. After all, one of the most important contributions of agent-based modelling is to explain complex patterns at the meso- and macro-level through simple micro-foundations and elaborate on mechanisms of emergence. Third, it is reasonable to test and implement a first-use case of data assimilation with rather simple ABMs. First we need to understand how to optimally integrate the two research strands, and evaluate the performance of the merger, before advancing to more complex models. Of course, it is expected that the trend towards large and high-dimensional agent-based models, including temporal and spatial granularity in the wealth distribution, continues. In the long-term, our approach might offer more advantages the larger a model becomes. Models which approach 'digital twin' character, that is they are more or less direct replicas of entire economies, and are fed with a lot of uncertain high-frequency data, are potentially suited best for data-assimilation-based control.

The remainder of this paper is structured as follows: In Section 2, we review related literature. In Section 3, we introduce our methods including the agent-based models applied, as well as the Ensemble Kalman Filter. In Section 4, we demonstrate results and lastly, in Section 5, we discuss our findings.

2. Related literature

2.1. Agent-based models of wealth dynamics

Besides ABMs, there are many economic modelling approaches that deal with inequality and wealth dynamics in particular. This includes various types of large computable general equilibrium models. Among new work is Blanchet et al. (2022), Blanchet and Martinez-Toledano (2023) who apply stochastic differential equations to capture the evolution of wealth. However, ultimately our goal is to integrate an economics ABM with a data assimilation approach and therefore we do not discuss in detail the merit of other modelling approaches compared to agent-based models.

In any case, agent-based modelling is a tool that is increasingly used in economics (Axtell and Farmer, 2022) and so it is for the analysis of economic inequality in particular. Recent agent-based studies of income inequality include research into the impact of inequality on the development of modern financial systems (Botta et al., 2021), how inequality impacts innovation and economic growth (Caiani et al., 2019), and the degree to which trickle-up growth — the idea that redistributing income towards poorer households stimulates overall economic growth — may reduce inequalities (Palagi et al., 2023). A popular class of agent-based inequality models, that is particularly relevant here, emerges from the econophysics literature. Dragulescu and Yakovenko (2000), Drăgulescu and Yakovenko (2001) and Yakovenko and Rosser (2009) show that the exponential-like distribution of income and wealth in the USA can originate from a simple model of agent-exchange. Here agents randomly bump into each other and exchange a unit of 'wealth', much like molecules do exchange energy in an ideal gas. Hence this strand of literature also has been termed 'Statistical mechanics of money'. The sociologist Angle (1986) also developed a similar model 15 years earlier. It has however been pointed out that this class of models fails to represent any meaningful economic ontology, in the sense that they do not consider any actual market mechanisms and so forth, and also fail to reproduce the essential fat tails of the wealth distribution (Lux, 2005). We do borrow heavily from this class of models in our own attempt to model wealth inequality since we also implement a simple exchange mechanism in our second model. But our approach for this newly developed model still diverges in an important way — the exchange that two agents have is constrained not only by their own quantity of wealth but also by their position in a social network which immediately gives rise to a more realistic distribution. The interplay between social network structure and the historical as well as contemporary evolution of inequality is an active area of research (Chiang, 2015; Mattison et al., 2016). Hence it seems reasonable to assume such a model.

A model we use as a starting point to describe the evolution of wealth and to integrate with data assimilation is by Vallejos et al. (2018). This model is called agent-based by the authors although microsimulation might be a better description. We argue so because the model has little to no agent-interaction and also the agents actually do not apply any 'agency' in the sense that they execute actions in response to some environment. Instead, the core-mechanism of this model is an aggregate economic growth function that, in discrete time, generates a certain 'pot of resources' that is distributed among the agents. So why even consider it? The model does fit surprisingly well to the historical time-series of wealth inequality, including its Pareto-tail, and exhibits some computational features of ABMs such as discrete individuals. Therefore it serves well as first test-bed to apply data-assimilation to.

2.2. Data assimilation and economics

Data assimilation (DA) is essentially applied control theory. It is a set of algorithms that optimally combine observations with model outputs to predict the future state of any given system. It has been largely developed in the atmospheric sciences and has been hugely successful in making weather predictions better (Kalnay, 2003) and also has manifold applications in control engineering.

It is key however to understand that data assimilation, and especially Kalman Filter (KF) approaches, are not new to economics. On the contrary the first applications date back several decades (Vishwakarma, 1970; Athans, 1974), and many more such applications have emerged over the following decades (Schneider, 1988; Pasricha, 2006; Thamae et al., 2015; Inglesi-Lotz, 2011; Munguia et al., 2019). Now there is also first work applying certain algorithms like particle filters to estimating agent-based models in economics - particularly with respect to latent parameter estimation and micro-level estimates (Lux, 2018; Monti et al., 2023).

Outside of economics, there is some preliminary work that integrates data assimilation with ABMs. These applications focus on urban dynamics and indoor or outdoor pedestrian dynamics as well as transport applications in particular (Wang and Hu, 2015; Ward et al., 2016; Malleson et al., 2020; Clay et al., 2021; Ternes et al., 2022; Hu, 2022; Tang and Malleson, 2022). Urban systems evolve in a fast-paced manner but are monitored by many different sensors such as cameras, air pollution sensors, traffic counters, etc. To make predictions about the system — such as crowd density, traffic congestion or pollution levels — in 'real-time', data assimilation algorithms are designed to optimally combine model predictions with incoming sensor data to use "all the available information, to determine as accurately as possible the state of [the system]" (Talagrand, 1991). In economics, there is an argument to be made, that as models grow larger, with more parameters, a wider range of potential system states, and especially an expanding volume of real-time economic data, the case for data assimilation as a method for optimal control between models and data is stronger than ever. This is where our work begins.

3. Methods

Here we described the different methodological components that make up the overall process. In Section 3.1, we provide a brief overview of the data we employ. In sections 3.2 and 3.3 we describe the agent-based models that we use, in Section 3.4, we elaborate on model calibration to historical data, in Section 3.5 we explain the Ensemble Kalman Filter (EnKF).

3.1. Data

We employ a monthly time-series of the wealth distribution constructed by Blanchet et al. (2022) as our observational data. With respect to the evolution of the wealth distribution, the monthly interval can be considered 'high-frequency' and a step towards 'real-time' observation since otherwise only yearly intervals had been available. The authors estimated the time-series based on individual tax accounts and through rescaling of monthly macroeconomic aggregates, such as GDP. This data is essentially a granular extension of the data that Piketty described in "Capital in the 21st century". The data features several wealth groups and we focus on the following disjoint groups: the top 1%, the next 9%, the middle 40% (50th to 90th percentile) and the bottom 50%. While Blanchet et al. (2022) report respectable accuracy for those groups compared to other estimation methods (meaning the standard deviation of a fit against another method for lesser frequency data is usually within 5% of the mean wealth of a group), the data itself is constructed based on various modelling assumptions and hence can be regarded as inherently uncertain. For example Auten and Splinter (2019) arrive at substantially different estimates of income inequality compared to Blanchet et al. (2022) but they did not yet account for wealth inequality as well. In any case, in principle large uncertainties are no problem for us, as this exactly why we work on this in the first place and we will represent data uncertainties further down in our hybrid-ABM-data-assimilation exercise. For simplicity we might regard the estimation of Blanchet et al. (2022) also as a "ground-truth", that is the true values of wealth inequality, which is necessary in order to know how well our model-data-assimilation approximates reality. In the following, we mostly regard the observations to be relatively certain, assuming the uncertainty of each observation for each wealth group to be defined via one standard deviation being 5% of the mean as a normally distributed random variable, in line with errors reported in Blanchet et al. (2022). This means if the average wealth estimate for one wealth group is 1000 dollars, it constitutes the average of a normal distribution with 50 dollars as one standard deviation to generate data ensembles for the Ensemble Kalman Filter. For experiment one, we assume substantially higher observation uncertainty because we investigate a longer time-series and this helps preserving uncertainty and variety in the model ensemble (discussed in detail in Section 4.1).

3.2. Agent-based model 1

We adapt the model created by Vallejos et al. (2018). For full details of the model parameters and calibration the interested reader can refer to the original study. Although simple, the model is able to reproduce the major wealth distribution trends as shown extensively in Vallejos et al. (2018).

At an aggregate level, any economy might be characterized by its measure of output or wealth and its growth rate. This also holds for individuals. That is, if we denote wealth as x and the constant growth rate per unit of time as g, a two-dimensional time-discrete model of any i agent's wealth can be given by (1) and (2):

$$x_{i,t} = x_{i,t-1} + x_{i,t-1} \cdot g_{i,t-1} \tag{1}$$

$$g_{i,t} = f(x_{i,t-1}) \tag{2}$$

where, following Vallejos et al. (2018), the growth rate of wealth is a function of wealth itself. Wealth returns are positively correlated with wealth (Fagereng et al., 2020; Palagi et al., 2023) and indeed the model fits well to the empirical data. Here $f(x_{i,t})$ is determined via the probability that an agent receives an additional unit of wealth. This is the defining feature of the entire model. Agents' trajectories are solely determined by a probability to receive more wealth. Vallejos et al. (2018) define this probability ψ as the share of wealth that an agent owns out of total wealth:

$$\psi(i,t) = \frac{p(x)_{i,t}}{\sum_{n=1}^{k} p(x)_{i,t}} \tag{3}$$

Note that $p(x) = x^{\beta}$. The parameter β is an exponent which determines how "social power" to receive additional wealth scales with wealth itself. The final major model assumption is that the economy grows exponentially; $\Omega(t) = \Omega(0)e^{\lambda t}$ where Ω is the wealth of agents summed up and the growth per time step $\Omega(t) - \Omega(t-1)$ is divided into n different discrete pieces of the 'growth-pie' and allocated to the agents via the above outlined probability rule.

3.3. Agent-based model 2

The second model, we developed from scratch. It is likewise an agent-based model of wealth distribution, although this new model incorporates agent interactions via a network. As the ability to model interactions is one of the key features of the methodology (Gilbert, 2008; Axtell and Farmer, 2022), we use this model because it is more characteristic of ABM in general. Although its ability to predict trajectories is slightly lower than the model by Vallejos et al. (2018), this is no detriment as our principal goal is to use the EnKF in combination with an imperfect ABM whilst still producing output that is close enough to reality to allow for an appropriate case study. There is a relatively rich econophysics literature studying the wealth distribution originating from network dynamics (Souma et al., 2003; Ichinomiya, 2012; Hu et al., 2008; Lee and Kim, 2007; Di Matteo et al., 2003), including work that shows that scale-free networks lead to power-law distributions — a fact we utilize here. Yet it is important to note that our goal is not to deliver a fundamental contribution to economic theory but to conduct a second test for the EnKF. Hence the model assumptions and dynamics may be considered a useful toy model rather than sophisticated economic dynamics.

The agents in the model are located on a Barabási–Albert (BA) network. Thus the model is generated via a preferential attachment algorithm and produces a power-law distribution of node degree. The power-law distribution is commonly expressed in the form $p(k) \sim k^{-\gamma}$, where $p(k)$ is the probability of a node in the network having degree k, and γ is the exponent of the power law distribution. In theory $\gamma = 3$ (Albert and Barabási, 2002), but given our small network sizes we find $\gamma \approx 2$ when measured empirically from the generated network. The only free parameter in a BA-graph is m which determines the minimum node degree. We set $m = 2$.

The assumption behind the network is that economic exchange between agents is constrained by social structure. Abstracting beyond households and firms, in reality, not everyone trades with everyone and some individuals have much more access to trade and resources than others. In the following, we then initialize the agents according to a heavy-tailed wealth distribution and associate the wealthiest agents with the highest node degree. Therefore, the richest agent has the highest node degree, and the poorest agent the least.

We assume that every agent's wealth $x_{i,t}$ grows at every time step by a global fixed rate, so $x_{i,t} = x_{i,t-1} * (1 + g)$. The core mechanism for wealth distribution dynamics, and hence for effective growth rates per agent, is trade however. Each agent trades every time step with all its network neighbours. Therefore, there are two trades between agents at each time step. If agent i initiates trade with j, symmetrically j initiates trade with i. The agents have a willingness-to-risk (WTR), which is the fraction of their own wealth they are willing to risk losing in a trade and which we denote by η. This parameter is drawn from a uniform distribution in the interval (0,u), where u is set to a maximum of one-tenth (to not generate trades too large). The trade amount between two agents i and j, denoted γ for fraction, is then the minimum of both quantities offered, so:

$$\gamma(i,j,t) = \min(x_{i,t}\eta_i, x_{j,t}\eta_j) \tag{4}$$

This is a reasonable first model assuming rational agents that offer only as much as they can afford but in line with diverse risk preferences. Moreover, the chance for any agent i to win the trade is determined by its network position, that is its node degree D_i subject to an exponent c. The latter lets us control how strong the difference in likelihood between two agents is due to their

Table 1
Model 1 key parameter names, values and references (where applicable).

Parameter	Symbol	Unit	Value	Reference
Growth rate	–	$\%\text{yr}^{-1}$	3	Blanchet et al. (2022)
Start year	–	yr	1990	–
Population size	N	–	100	–
‘Social power’ parameter	β	–	1.3	Vallejos et al. (2018)

different network position but we set it generally to one. A more central agent is more likely to win, representing a kind of ‘social power’:

$$P_{win}(i,t) = \left(\frac{D_i}{D_i + D_j}\right)^c + \epsilon_t \tag{5}$$

Here ϵ_t is an error term drawn from a truncated normal distribution.

Lastly, and importantly, agents adjust their WTR based on whether they win or lose or the trade. If they lose, they lower their WTR, if they win they increase it. Therefore the agents are adaptive in nature. We call this parameter adaptive sensitivity or s (see Table 2).

This model converges to a steady state or equilibrium because over time when $t \to \infty$, the WTR of an agent who has on average a lower node degree than its neighbours goes towards zero, while the WTR of agents that have on average higher node degree than their neighbours tends towards the maximum value. At some point, trading halts because certain agents are no longer willing to engage. The equilibrium prediction for wealth distribution is plausible, as it qualitatively aligns with U.S. wealth data: the top 10% of the population hold well over half of all wealth, while the bottom 50% hold virtually none. If the adaptive WTR is taken out of the model, in the long-run all wealth ends up with the top 10%, which is less plausible. Section 3.4 illustrates the accuracy of the model in detail.

3.4. Model calibration and benchmarking

First, we initiate the models with the appropriate distribution of wealth at time point $t = 0$. We choose the period of January 1990 to December of 2022 as our benchmark period to test the models and the EnKF. There are several interesting crisis moments that occur during this period (e.g. financial crisis 07/08 and Covid pandemic) that provide a useful test for the methods, but at the same time the data includes many “early” years that are likely more reliable than newer estimates that arose during the pandemic and post-pandemic period and hence more appropriately apply as “ground truth”. Initially, we fit the agent wealth distribution to the empirical distribution of 1990. We do so based on a distributional model from Vallejos et al. (2018) because they achieve a good fit to the US wealth distribution (Refer to Eq. (8) in their paper). This probability distribution is a weighted average of an exponential distribution and a pareto distribuion and its quantile function reads as follows.

$$Q(p) = -T \cdot (1 - \pi(p)) \cdot \ln\left(\frac{1-p}{c}\right) + \pi(p) \cdot \omega \cdot (1-p)^{-\frac{1}{\alpha}} \tag{6}$$

where $Q(p)$ is the wealth at a certain percentile, T is a parameter, $\pi(p)$ is the weights parameter dependent on the percentile, p is the given percentiles, c is a constant (in our calibration just set to 1), ω is the minimum wealth parameter, α is the shape parameter of the Pareto distribution.

We also set the weights as they do in Vallejos et al. (2018) via a piece-wise function - although we choose slightly different weights. This is because Vallejos et al. (2018) employ a different segmentation of the data. For example, they lump all the bottom 90% together which simplifies the exercise as the distinction between the bottom 50% (lower wealth classes) and the next 40% (middle class) entails additional structure. In our calibration, up to the 40th percentile only the exponential distribution applies, then between the 40th and the 90th percentile a weighted average and beyond the 90th percentile only the Pareto-distribution. Lastly, we apply a scaling coefficient to scale the sample to the appropriate expected mean value which corresponds to the data. For details there is an accompanying Jupyter notebook in the corresponding Github repository.

We initialize 100 agents by sampling from this distribution for each simulation, which is a rather small scale but we aim for efficiency to test the EnKF algorithm. A small sample size of agents guarantees also a sufficient variance across the ensemble. We apply inverse-transform sampling to sample from the defined quantile function.

Both models simulate the compound annual growth rate in the USA from 1976–2023 which is ∼3% according to the data by Blanchet et al. (2022) as key parameter for driving dynamics. We list all the key parameter specifications in Tables 1 and 2.

In all computational experiments, we employ a mean absolute error metric (MAE), defined in Eq. (7) to evaluate the performance of the models as well as that of the EnKF with respect to the ‘ground truth’. The ground truth is defined as the raw data given by Blanchet et al. (2022) and assuming no uncertainty in the data. Note that the EnKF creates a number of distinct instances of the model and executes them simultaneously (full details are forthcoming in Section 3.5). Together these model instances are called an ‘ensemble’ and an ‘ensemble member’ refers to one of those instances. Therefore in Eq. (7), m_{ijt} is the prediction of the ith ensemble

Table 2
Model 2 key parameter names, values and references (where applicable).

Parameter	Symbol	Unit	Value	Reference
Growth rate	–	$\%yr^{-1}$	3	Blanchet et al. (2022)
Start year	–	yr	1990	–
Population size	N	–	100	–
Willingness-to-risk (WTR)	η	–	$\in (0,0.1)$	–
Adaptive sensitivity	s	$\%t^{-1}$	2	–

member for the jth wealth group at time point t, and d_{jt} is the ground truth for the jth wealth group. We take the absolute value difference and sum over the wealth groups and then over the ensemble members and finally average across ensemble members.

$$MAE_t = \left(\frac{1}{n}\right) \sum_{i=1}^{n} \sum_{j=1}^{4} \left| m_{ijt} - d_{jt} \right| \tag{7}$$

Fig. 2 shows the calibration results. In Fig. 2 panel (A) and (B) both models are depicted with one singular run against the data to illustrate their archetypal behaviour. We see that Model one produces more "linear" trajectories while Model two produces non linear trajectories but tending towards the above mentioned equilibrium. In reality, at times adults from the bottom 50% may exhibit negative wealth which in both models on their own does not occur. Keep in mind that all the wealth data is real wealth per adult person. In Fig. 2 panels (C) and (D) we depict the uncertainty with confidence intervals across 100 simulation runs. In both models, the top 1% group exhibits the most variation across simulations. The high variability in this wealth group is fitting, given the uncertainty surrounding the wealth of the very wealthy. Finally, Fig. 2 panel (E) depicts the error as defined in Eq. (7). Model 2 exhibits higher error in the long run but its error stabilizes, while the Model 1 error increases slightly over time. For our purposes, testing the EnKF with agent-based modelling in economics, it is of advantage to work with both models so that we can find out whether the EnKF works under both circumstances equally well.

3.5. Ensemble Kalman filter

In the following section, we outline the Ensemble Kalman Filter (EnKF) and its application to our agent-based wealth distribution model. Note that a simple example is included at the end of the section to aid understanding of the method. The Ensemble Kalman Filter is a Monte-Carlo variation of the Kalman Filter (Kalman, 1960). The original Kalman Filter is a sequential data assimilation approach which is applied to a model of a system which has an associated set of state variables ($\mathbf{x}$) and forecasting process; the forecasting process describes the way in which the state variables (and associated covariances ($\mathbf{Q}$)) are evolved over time. When observations of the system ($\mathbf{d}$) are provided, the Kalman Filter updates the model state based on the relative covariances of the observations and the model state, providing an updated estimate of the model state, $\hat{\mathbf{x}}$:

$$\hat{\mathbf{x}} = \mathbf{x} + \mathbf{K}\left(\mathbf{d} - \mathbf{H}\mathbf{x}\right), \tag{8}$$

and an updated model covariance matrix, $\hat{\mathbf{Q}}$:

$$\hat{\mathbf{Q}} = \left(\mathbf{I} - \mathbf{K}\mathbf{H}\right)\mathbf{Q}. \tag{9}$$

In this context, $\mathbf{K}$ is the Kalman Gain matrix and $\mathbf{H}$ is the observations operator, both of which shall be expanded upon in the below explanation of the Ensemble Kalman Filter.

The Kalman Filter makes two key assumptions: that the errors in model estimates and observations are described by a Gaussian distribution, and that the dynamics of the system are linear. As such, its performance may suffer in contexts where these criteria are not fulfilled. Furthermore, it requires that the covariance matrix, $\mathbf{Q}$, be maintained — both when forecasting the model state over time and when updating the model state based on observations. This process can become computationally expensive for high dimensional models (Evensen, 2009).

The EnKF is particularly useful in situations where the underlying model is complex and non-linear (Evensen, 2003) — as is the case with agent-based models. Compared to the original Kalman Filter, the EnKF has several advantages:

Computational Efficiency: the Ensemble Kalman Filter uses an ensemble of state estimates that can be propagated independently in parallel, providing the ability to derive covariance matrices in a more computationally efficient manner for large state spaces;

Non-linearity: it handles non-linearity more effectively as the state covariance matrix no longer needs to be forecast by applying a linear operator, but instead can simply be generated as a sampling covariance.

The basic idea behind the EnKF is to represent a system and its associated uncertainty using a collection (or "ensemble") of possible state estimates, rather than just a single estimate as would be the case with the original Kalman Filter. These ensemble members represent different possible evolutionary paths of the system, based on different assumptions or uncertainties in the model. At each time step, the EnKF uses the model to predict the state of the system based on the current ensemble of state estimates

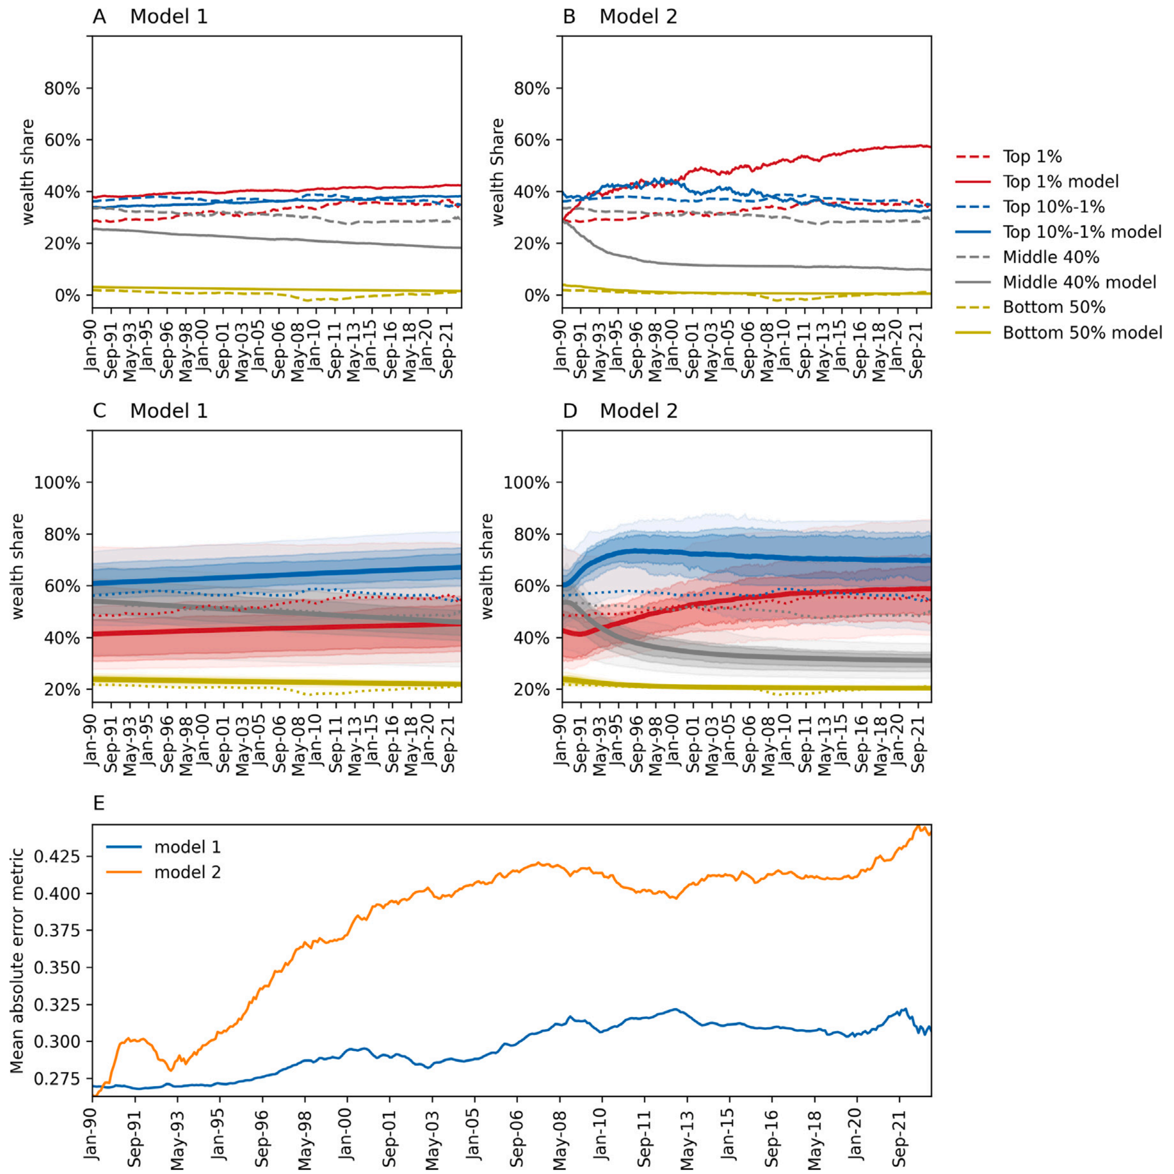

Fig. 2. Calibration results: illustrating the error of models 1 and 2 compared to the real data. Panels (A–D) depict the wealth share of the different economic groups (top 1%, top 10%–1%, middle 40%, bottom 50%) over time for the models' predictions and the real wealth data. Panels (A) and (B) present the archetypal behaviour of a single model run. Panels (C) and (D) show the mean prediction and uncertainty across 100 model runs. Panel (E) depicts the Mean Absolute Error (MAE) of the two models as per Eq. (7).

(termed the 'predict' step). At time steps when observations of the system are provided, it then compares these predictions to actual observations of the system and adjusts the ensemble to better estimate the system state (the 'update' step). The key to the effectiveness of the EnKF (and data assimilation approaches more generally) is that it uses information from both the model and the observations to provide a best estimate of a system. This allows it to handle situations where the model is imperfect or incomplete, and where there may be substantial uncertainties in the observations.

The ensemble of system state estimates, $\mathbf{X}$, that consists of a number of individual realizations of the model, $\mathbf{x}$, is defined as:

$$\mathbf{X} = \left[\mathbf{x}_1, \dots, \mathbf{x}_N\right] = \left[\mathbf{x}_i\right], \quad \forall i \in \{1, 2, \dots, N\}, \tag{10}$$

where N is the number of ensemble member models. The mean state vector, $\bar{\mathbf{x}}$, can be found by averaging over the ensemble:

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{x}_i. \tag{11}$$

Similarly, we have a collection of data vectors, **D**, which are considered to be "nearly-true":

$$\mathbf{D} = \left[\mathbf{d}_1, \ldots, \mathbf{d}_N\right] = \left[\mathbf{d}_i\right], \quad \forall i \in \{1, 2, \ldots, N\}. \tag{12}$$

These "nearly-true" data vectors are generated using a simple stochastic rule; each data vector consists of the sum of the original observation **d**, and a random vector, ϵ_i:

$$\mathbf{d}_i = \mathbf{d} + \epsilon_i, \quad \forall i \in \{1, 2, \ldots, N\}. \tag{13}$$

The random vector is drawn from an unbiased normal distribution:

$$\epsilon \sim \mathcal{N}(0, \mathbf{R}), \tag{14}$$

where **R** is the data covariance matrix which is defined by the uncertainty in the observations.

Given the above framework, the data assimilation is made up of the predict-update cycle (outlined in Fig. 3), with the calculation of an updated state ensemble, $\hat{\mathbf{X}}$, being undertaken on the basis of the following equation:

$$\hat{\mathbf{X}} = \underbrace{\mathbf{X}}_{\substack{\text{agent}\\ \text{states}}} + \underbrace{\mathbf{K}}_{\text{weight}} \underbrace{(\mathbf{D} - \mathbf{H}\mathbf{X})}_{\text{perturbation}}, \tag{15}$$

where **H** is the observation operator and **K** is the Kalman gain matrix. The role of the observation operator is to transform the state vectors between the form in which we store state variables (in our case, wealth per agent) and the form in which they are represented in observations which is a macroeconomic time-series (namely wealth per wealth group). An important clarification is that based on (15) we only update the state of the model, that is the agent wealth variable, but not parameters determining the model state. In principle one can also use data assimilation to learn the underlying parameters but for our case study we focus on the state variables to keep it relatively simple.

Note here that we do not update our state covariance matrix, instead generating a sample covariance matrix based on the state ensemble, **C**:

$$\mathbf{C} = \frac{1}{N-1} \sum_{i=1}^{N} \left(\mathbf{x}_i - \bar{\mathbf{x}}\right) \left(\mathbf{x}_i - \bar{\mathbf{x}}\right)^T.$$

The Kalman gain matrix, **K** is given by

$$\mathbf{K} = \mathbf{C}\mathbf{H}^T \left(\mathbf{H}\mathbf{C}\mathbf{H}^T + \mathbf{R}\right)^{-1}, \tag{16}$$

and **R** is the observation covariance. We can consider (**D** − **HX**) in Eq. (15) to be the proposed perturbation to the ensemble states, and the Kalman gain matrix, **K**, to be the weight given to this perturbation (just as in a standard Kalman Filter) based on the relative uncertainties in the state ensemble and the observations. When the uncertainty in the observations is low in comparison to the uncertainty in the model state, the gain increases, and consequently the model state receives a larger perturbation from the provided data; conversely, when the uncertainty in the observations is high in comparison to the uncertainty in the model state, the gain decreases, and consequently the model state receives a smaller perturbation from the provided data. Fig. 3 provides an overview of the flow of the combined algorithm using the EnKF with the respective agent-based model.

Before outlining the results, it is useful to illustrate briefly how the agent states, observations and the observation operator are organized in matrix form. Assume a wealth agent-based model with only three agents, one rich one and two poor ones and three ensemble simulations. The exact mechanics of the model do not matter because we are interested in understanding how the EnKF alters the output and the state of the model. Moreover let the observations come in two groups, rich and poor. We also assume particular values, so the numbers in the matrices are the wealth if not specified otherwise. Then the entire ensemble of model states and the observation can be depicted as follows. Here $\mathbf{x}_i$ are the ensemble members and the dashed lines indicate the agent groups that correspond to specific observation wealth groups.

$$\begin{array}{l} \\ \text{Agent}_1 \\ \text{Agent}_2 \\ \text{Agent}_3 \end{array} \begin{array}{c} \begin{array}{ccc} \mathbf{x}_1 & \mathbf{x}_2 & \mathbf{x}_3 \end{array} \\ \begin{pmatrix} 150 & 170 & 130 \\ 10 & 15 & 5 \\ 5 & 5 & 5 \end{pmatrix} \end{array} \tag{17}$$

$$\text{Obs.}_{\cdot \mathrm{i}} = \mathbf{d}_\mathbf{i} = \begin{pmatrix} \textit{rich} \\ \textit{poor} \end{pmatrix} = \begin{pmatrix} 150 \\ 11 \end{pmatrix} \tag{18}$$

Now the observation operator, **H**, compresses the 3 × 3 matrix of agent states to a 2 × 3 matrix so it matches the observation dimensions of some $\mathbf{d}_\mathbf{i}$ and **D**. If there is more than one agent within a group, as in the case of the poor agents, then **H** averages over the state variable. In our case this operator effectively bridges the micro-economic level and the macro-economic level via aggregation.

$$\mathbf{H} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0.5 & 0.5 \end{pmatrix} \tag{19}$$

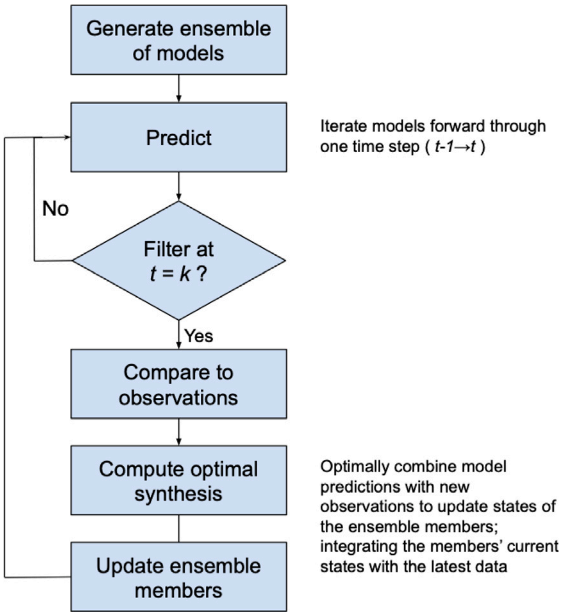

Fig. 3. Ensemble Kalman Filter algorithm overview. Note that t is the model timestep and k indicates whether t is an *update* step (i.e. a timestep when the EnKF updates the model state in response to new data).

A key clarification regarding model two is that we introduce an additional step in Eq. (15) to prevent agent states from updating to wealth values below zero. This adjustment is necessary to avoid divergence to negative infinity in model two. Although other options could achieve a similar outcome — such as modifying the model to exclude agents with negative wealth from all growth and trade processes — we chose this approach to highlight unresolved issues with the EnKF in agent-based wealth models. This choice emphasizes the need for alignment between model mechanisms and the filter's scope of application to ensure effective integration in future research. We include however a variation of this analysis, where we altered the model instead, and without the intervention in the filter step, in the supplementary information and it demonstrates that the approaches are largely equivalent in outcome.

4. Results

4.1. Experiment 1: Filter vs. no filter

The first experiment compares the EnKF-supported models with the baseline performance depicted in Fig. 2. The model parameters are set as in Tables 1 and 2. For the uncertainty in the observations we make pragmatic assumptions as it is not clear how uncertain the data are. In Blanchet et al. (2022) there is an analysis of predicted vs. actual growth rates when testing their method for estimating real-time wealth inequality against established "slower" wealth accounting methods. They report generally low standard deviations of between 4% and 6% of the mean value depending on the exact wealth group considered. For the sake of simplicity, we may assume a standard deviation of 5% of the mean value for observations for each wealth group which is well in line with their estimates. We do test varying uncertainty assumptions in Section 4.5, but in this specific experiment we retain a much higher uncertainty in the data (standard deviation of 50%) so that the model ensemble variance does not collapse to a low value - because if the data is assumed too certain, the EnKF does not retain much variance in the model ensemble .

For this experiment, we distinguish between two sub-parts (i) testing the behaviour of the EnKF in conjunction with both models over time and (ii) analyse in detail how the Filter update alters the agent-states, that is the distribution of wealth at the agent-level.

4.1.1. Time series perspective

We apply the filter every 20th time step using 100 ensemble members and models with 100 agents. We do not use a smaller number of ensemble members for this experiment because as we will find later in, Section 4.3, that this influences the reliability and stability of the results (illustrated in Fig. 7). The filter frequency is set to a moderate level, allowing sufficient time for individual simulations to evolve independently of the filter, thereby preserving variance across the ensemble. Fig. 4 outlines the results. Panel E demonstrates that, as expected, the error between ground truth and model is substantially reduced with the use of the EnKF. Panels A and B demonstrate, for illustrative purposes, single simulations of models 1 and 2, chosen at random, and how the EnKF affects them. Panels C and D illustrates the same but for the ensemble of all 100 members (solid lines show the mean of all ensembles). Considering panels C and D, it is apparent that the first correction step, around Sep-91, makes a substantial change to the trajectory of both models. We can observe the following:

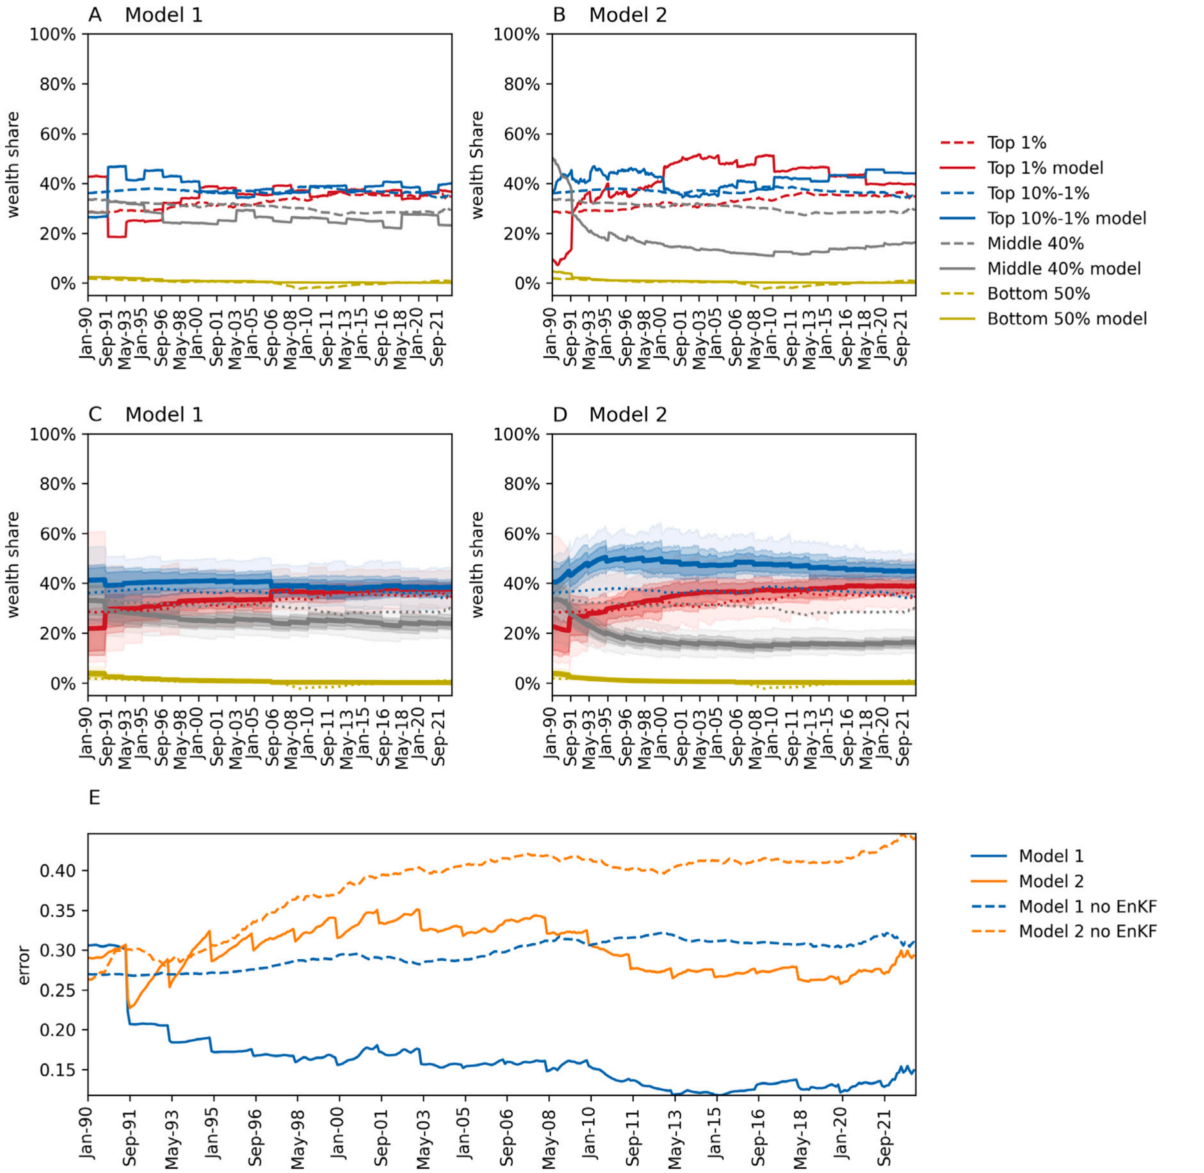

Fig. 4. Results of Experiment 1 Part A: illustrating the error of models 1 and 2 under Ensemble Kalman Filter (EnKF) optimization with 100 ensemble members compared to the real data. The filter is applied every 20 time steps (months). Panels (A–D) depict the wealth share of the different economic groups (top 1%, top 10%–1%, middle 40%, bottom 50%) over time. Panels (A) and (B) present the archetypal behaviour of a single model run, illustrating how the EnKF influences the model behaviour. Panels (C) and (D) show the mean EnKF prediction and uncertainty across all ensemble members. Panel (E) depicts the Mean Absolute Error (MAE) from Eq. (7) of the two models, with and without the EnKF. It is clear that when the EnKF is applied (once every 30 iterations) the model error decreases. Note that the exact trajectory of the error can vary by experiment instance.

- The bottom 50% are corrected down towards their actual true state value which is near 0% of wealth.
- The top 1% are corrected upwards to reflect a higher proportion of wealth ownership than the models predicted.
- The proportion of the top 10%–1% and middle 40% of agents are not corrected as well as other groups. The possible causes of this finding will be discussed in Section 5.

The important lesson from this first experiment is that, overall, the error over time remains substantially lower than when using no EnKF - around 20% to 50% for model 1 and around 20% to 40% depending on the time step (these numbers may vary somewhat when running another instance of the experiment).

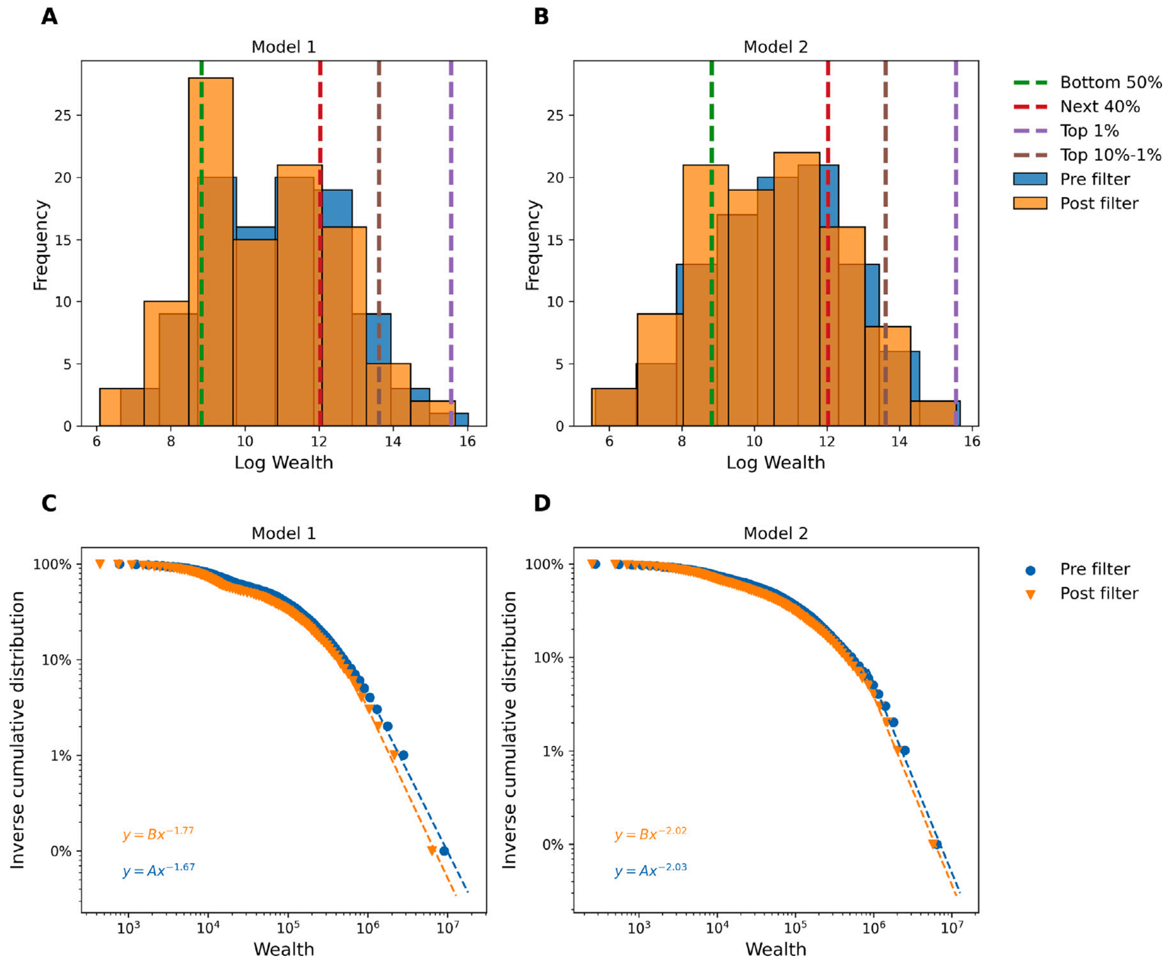

Fig. 5. Results of Experiment 1 Part B: illustrating how the agent states, that is the wealth per agent distribution, is updated during the first filter step in Experiment 1. Panel A and B show the mean average distributions across the ensemble pre filter (blue) and post filter (orange) as histograms including the corresponding observations at these time steps. Panel C and D illustrate the inverse cumulative average distributions and show that the nature of the distribution is not affected much. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

4.1.2. Agent-states perspective during an update step

In this section, we illustrate the distribution of wealth over all agents pre- and post update a filter step for the mean average of the ensemble. Specifically, we show in Fig. 5 the first update step in the time-series from Fig. 4 but for individual absolute wealth values in log terms instead of shares per wealth group. We find that the filter adjusts the distribution of agent states by shifting the distribution slightly to the left. The filter reduces the wealth values of the Bottom 50% and Top 10% of agents, resulting in an increased share of the top one percent, so the top agent represents a larger portion of overall wealth. For example, in Panel A and B of Fig. 5, it is shown that the post filter distribution (orange) exhibits more agents being close to the observation of the Bottom 50% than the pre filter distribution (blue). Furthermore, Panels C and D display the inverse cumulative average distributions before and after filtering, as is common in studies of wealth distributions. The overall shape of the distributions remains largely unchanged.

4.2. Experiment 2: Variation of filter window size

Recall that the filter ‘frequency’ refers to the number of model time steps (in this case months) between assimilation episodes. Plausibly we would expect that higher filter frequency should result in an overall better estimate of the true state of the system as the models are not required to predict for too long without any support from the data. Here we test this assumption by varying the filter frequency and calculating the total error for the entire model run over time, S, by adding a third summation term to Eq. (7):

$$S = \sum_{t=1}^{t} \left(\frac{1}{n}\right) \sum_{i=1}^{n} \sum_{j=1}^{4} \left| m_{ijt} - d_{jt} \right| . \tag{20}$$

The results are illustrated in Fig. 6.

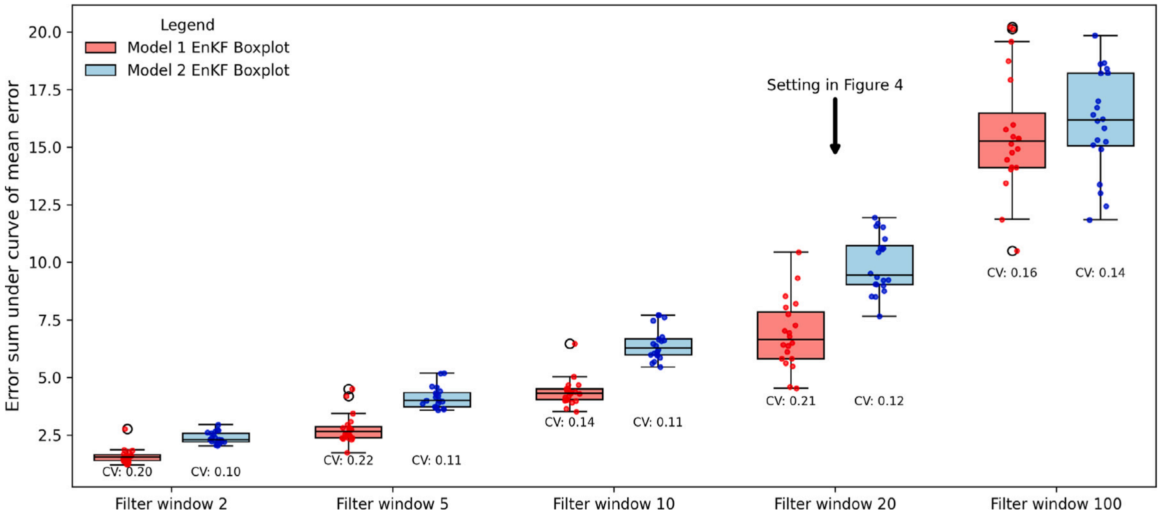

Fig. 6. Results of Experiment 2: the influence of the filter frequency on the overall error for models 1 and 2. Each filter frequency experiment is executed 20 times. As the window size (x axis) increases from 2 iterations to 100 iterations, the overall error of the EnKF run S (Eq. (20)) increases as expected. We test the robustness of this experiment with respect to other time periods in Supplementary Figure 2 and 3.

To keep this experiment computationally efficient, we limit the time horizon to the first four years, so 1990–1993, again with 100 agents, fix the ensemble size at 10 and repeat each filter frequency 20 times. We test filtering every 2nd up to every 100th time step. We do not test filtering at every time step; given that the model introduces a limited amount of variance across the ensemble over the course of a single time-step, the gain that is applied when updating is very small. As a consequence, regardless of whether the model ensemble is predicting in line with the ground truth, it is liable to diverge from the ground truth. This is because, as with every model, the model used is an imperfect representation of the system dynamics. One solution which could be used in future work if attempting to assimilate every time-step would be to use covariance inflation approaches (Li et al., 2009).

We quantify two further metrics for this experiment — the elasticity of the error with respect to the filter frequency and the coefficient of variation across experiment repetitions (one experiment parameter setting including all its repetition represents one of the box plots in Fig. 6). In economics, the elasticity of a variable with respect to another is $E_{x,y} = \frac{\%\Delta x}{\%\Delta y}$, where in our case x is the error and y is the filter frequency. The average error elasticity for both models is ∼0.5. This means that the error increases by 0.5% for every 1% increase in the filter window length. The coefficient of variation of each experiment, which is the ratio of the standard deviation to the mean, is plotted below the corresponding boxplot. Generally it is relatively stable across different filter frequencies.

4.3. *Experiment 3: Variation of ensemble size*

Next we isolate the effect of the ensemble size on the EnKF-model performance. Again we quantify the error using the time aggregated error S from Fig. 4 panel E, choose a short time horizon (1990–1993), and fix the filter frequency at 10 time steps. Each ensemble size is tested 20 times. We estimate the error elasticity with respect to ensemble size, which is approximately −0.1 for both models but with strongly diminishing returns — an ensemble of 30 simulations works nearly as well as one of 100. The main finding here is that smaller ensembles produce substantially less reliable outcomes. The coefficient of variation across experiment repetitions is roughly 0.3 or higher for both models when the ensemble size is 5 and only around 0.05 for both models when the ensemble size is 100. This means the EnKF should not be employed with ensemble sizes smaller than 30 at least with respect to our case study. In other words, larger ensemble sizes in model and observations gauge a more reliable estimator of the true state of the system. In terms of standard statistical terminology, ‘accuracy’ pertains to the proximity of the mean estimate to the true value, while ‘reliability’ relates to the consistency of estimated values around the mean.

4.4. *Experiment 4: Trade-off between filter window size and ensemble size*

In this experiment, we combine filter window size and ensemble size to quantify the trade-off between the two. Our variable of interest is again the error under the curve as in Fig. 4 panel E (Eq. (20)). Here we only repeat the experiment for each window-size-ensemble-size combination five times and we only test the behaviour over two years – all for computational efficiency. The main finding, depicted in Fig. 8, is that the size of the ensemble primarily influences the reliability of an estimate, as we know already from the previous experiment, with only a minor impact on its accuracy, but the filter frequency, that is s smaller filter window, improves accuracy substantially. An increase in ensemble size contributes to accuracy enhancement mostly only when the filter window is quite small (so filter frequency is high), because this preserves some variance across the ensemble despite filtering often.

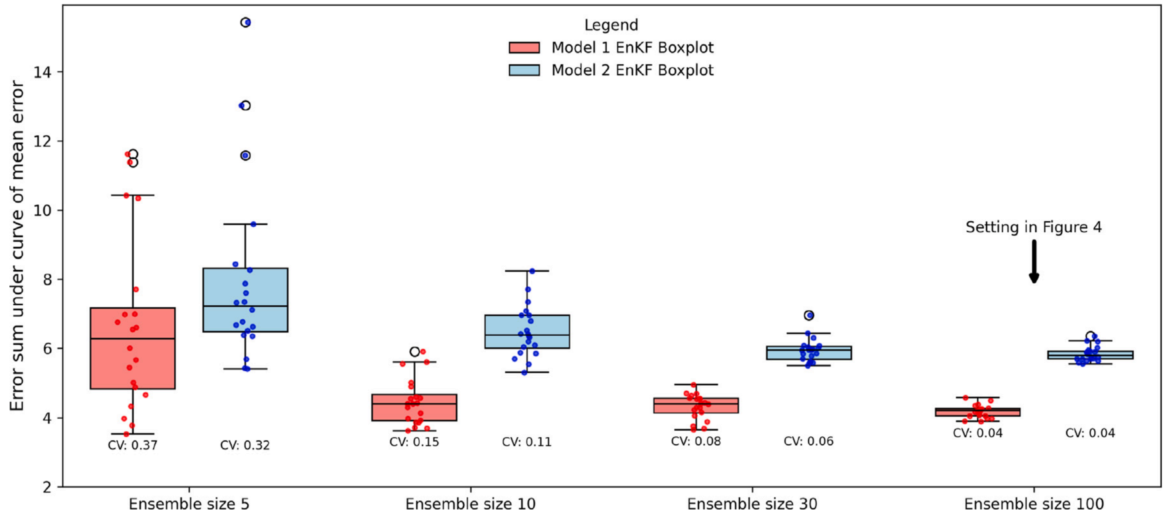

Fig. 7. Experiment 3: Exploring the influence of the EnKF ensemble size on the overall error (S, Eq. (20)). Each ensemble size experiment is executed 20 times. As the ensemble size (x axis) increases from 5 to 100 there is a decrease in error (y axis), as expected, but there are strongly diminishing returns. The error estimate becomes more reliable though.

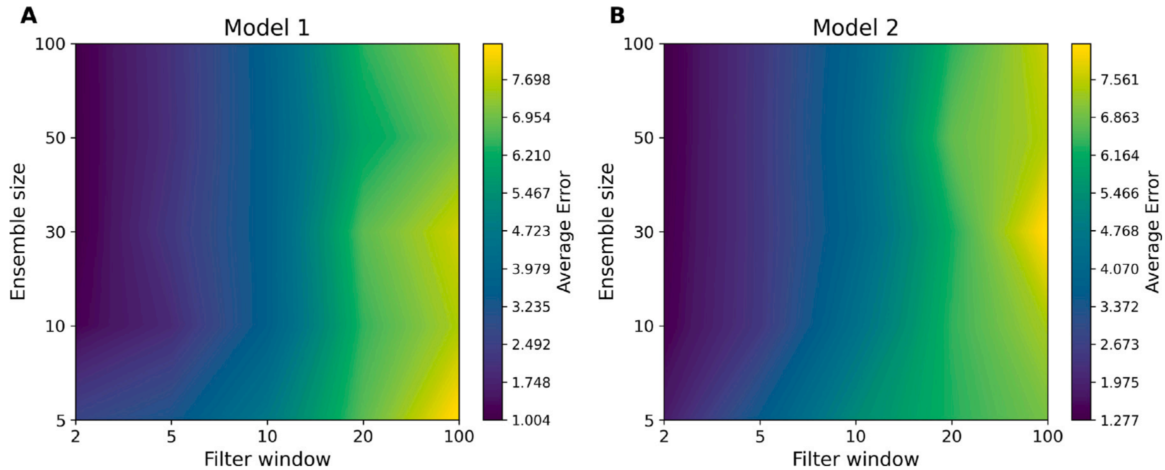

Fig. 8. Experiment 4: Comparing time aggregated errors (S) in response to changes in filter frequency (x axis) and ensemble size (y axis) for models 1 and 2. We display the mean error for each window-size-ensemble-size combination after 5 executions. It appears that the ensemble size has the more important impact on the reliability of the estimate; smaller filter windows result in a substantial decrease in error.

4.5. *Experiment 5: Variation of assumptions about uncertainty*

Experiment five studies the differing impacts that uncertainty in the model and uncertainty in the observations will have on the reliability of the filter. For artificially changing the observation uncertainty, we assume a range of different observation standard deviations. Our standard deviation was 5% of the mean of each wealth group for experiment two to four - in line with the uncertainty reported in Blanchet et al. (2022) - and for this experiment we vary this assumption from 5% to 100% (with more explanation on this observation uncertainty in Section 3.1). As models one and two are structured differently they require distinct adaptations to artificially vary their uncertainty. For model one we alter the "power" that an agent derives from having wealth by introducing a normally distributed term to the exponent that defines that "power" $p^{\beta+\mathcal{N}(0,a)}$. The term a represents the standard deviation of the normal distribution around the mean value of β. For model two, we simply vary the normally distributed error defined in Eq. (5). We vary the error for model one in terms of standard deviations from one tenth of a standard deviation to three standard deviations. For model two the error, because it alters a probability, can be meaningfully varied within the interval $[0, 1]$. We also define a new performance metric for this experiment because we are not interested in the error relative to the ground truth, but in how sensitive the EnKF is to changes in the model and observation uncertainty and whether these changes in uncertainty cause it to tend towards one rather than the other. Therefore we evaluate the posterior difference, T, between the mean model estimate $M_{j,i}$ for wealth

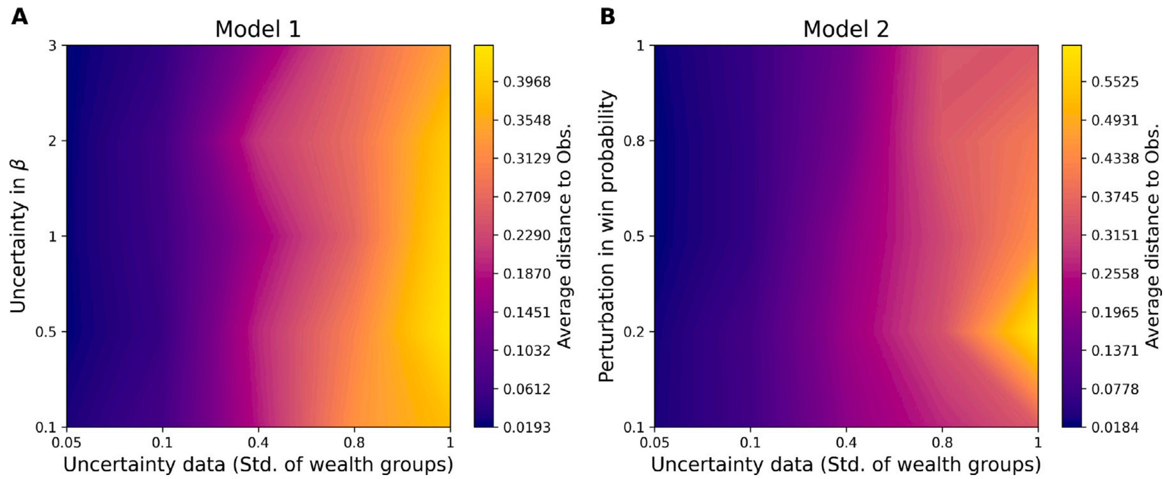

Fig. 9. Experiment 5: Trade-off between observation uncertainty (x axis) and model uncertainty (y axis) for models 1 and 2, measured using the posterior difference T; Eq. (21). As expected, we see that T decreases as observation uncertainty decreases. Each error distance measure is the average of ten experiment repetitions.

group j at time point i and the data ensemble average $D_{j,i}$. This is the average of **D** defined in Eq. (12). We then sum up across wealth groups and over time as depicted in Eq. (21):

$$T = \sum_t \sum_i \left| M_{j,i} - D_{j,i} \right| \quad (21)$$

The size of this posterior gauge should change in the same direction as uncertainty in the data (i.e. decrease as uncertainty in the data decreases and vice versa). Effectively we just measure the distance to the given observation which should increase as those become more uncertain. This is exactly what we find; Fig. 9 illustrates the results for both models. We run the experiment for two years and ten repetitions, a fixed filter frequency of ten and 20 ensemble members.

5. Discussion and conclusions

In summary, we have shown that the Ensemble Kalman Filter (EnKF) can be successfully combined with macro-economic agent-based models, specifically for the dynamics of wealth inequality. The EnKF consistently improves the fit to the ground truth state of the system, as characterized by the data provided by Blanchet et al. (2022) (Section 3.1). An important contribution of our work is that our hybrid implementation is tailored specifically to capture the dynamics of wealth distribution and, by extension, broader macroeconomic dynamics. This approach advances beyond other recent ABM-Bayesian-Update hybrid models like the one by Yang and Zhou (2022) who focus on reproducing the shape of the wealth distribution in a static setting.

Generally, while currently the benefit of making up-to-the-minute predictions is less significant for economics than for fields such as meteorology (Kalnay, 2003) or crowd simulation (Clay et al., 2021), the impact of creating a system to make “live” (Swarup and Mortveit, 2020) predictions via data assimilation (DA) and methods like the EnKF remains potentially transformative. For example, economists could create a simulation platform that made predictions about future conditions days, months, or years in advance. Using methods like the EnKF, the system could take up-to-date, diverse data sets from sources such as official government releases, social media data, new empirical research results, etc., and use the DA-ABM framework as a means to combine them and create accurate predictions using “all the available information” (Talagrand, 1991). Going further, through system state estimation, DA also holds a speculative potential for enhancing our *understanding* of economic systems, such as the wealth inequality dynamics explored here. In our case, for example, this could be by exploring the evolution of the Kalman Gain matrix (Eq. (16)) and by analysing the weight applied to the model state versus the observations, which in turn may provide insights into the conditions under which the model becomes less certain in its predictions and which of its features are of low confidence (e.g. which wealth groups). Moreover, there is also a substantial opportunity to update latent variables through data assimilation, as for example conducted in Monti et al. (2023). Latent state estimation is distinct from what we have done here because when DA methods like the EnKF correct only observation states, they only address the immediate, visible discrepancies between the model and reality. In contrast, by tuning latent states, the filter corrects the model’s internal assumptions and rules themselves. This approach means the filter can help the model “learn” more accurately from “real-time” data, which in turn can even reduce the need for frequent external corrections.

Moreover, incorporating latent variables into the DA framework offers the potential to learn the underlying dynamics in data-rich time periods (or, for a spatial model, data-rich *regions*) and apply this knowledge to data-poor contexts (Kalnay, 2003). This transferability could be explored through the use of out-of-sample forecasts, which would also have the advantage of further testing the model’s robustness. By generalizing learned patterns, the model could provide better forecasts even during non-crisis periods, or in new economic regions that it has not been calibrated on.

In our case study, an unexpected observation is that, occasionally, the EnKF makes the prediction worse (e.g. in Fig. 4 panel E). This can happen due to sampling error in the data ensemble, for instance, when the sampled observations diverges from the ground truth. The Ensemble Transform Kalman Filter (ETKF) is a variant of the EnKF that intends to overcome this problem (Zhou et al., 2019); future work could explore this approach as an alternative to the EnKF for economic models. In addition, the specific case of the U.S. wealth distribution explored here exhibits unusual features that can be problematic. For example, in 2007 and onwards there are negative wealth values for the bottom 50% and, simultaneously, large positive values for the very wealthy groups. This can make it challenging for the EnKF to find optimally balanced update weights for each group, particularly when group sizes are small. For example, in our case, a single agent represents the top 1%, while in the observed data — used to calculate the 'error' that the EnKF adjusts for — the top 1% is one of four groups, a distinct dimensionality.

We tested two dynamic models simultaneously for a more general performance evaluation of our proposed method. While both models in an EnKF-ABM setting act relatively similarly, Model Two with EnKF performs somewhat worse. This is because the steady-state of the model is further away from the ground-truth data than the average model run in Model One. Moreover, this steady-state is attracting. If the model is corrected towards the ground truth by the EnKF, there is a stronger tendency for the model to evolve towards its characteristic steady-state again. This is an important lesson for further studies combining economic models and data assimilation. Most classical economic models exhibit "transitional dynamics" towards a steady state (Acemoglu, 2008). While many agent-based models are constructed for the purpose of increased complexity and instability when compared to classical equilibrium models, they still might exhibit long-run steady states or even emergent, unforeseen, attractors. Our results suggest that if those steady states are rapidly converging and perhaps not exactly matching some more "varied" time series, it might be more difficult for data assimilation to learn the true state of a system. The fact that the ground truth state that we track is four dimensional (because there are four wealth groups), with possibly diverging trajectories, also amplifies this problem. For this reason, it is important for the EnKF to intervene particularly early on with respect to Model Two. If the data assimilation window is too wide, and the model has already diverged far from the ground truth, it is more difficult for the model-EnKF hybrid to generate a low-error performance over time.

In addition, there are many practical and theoretical implementation challenges for agent-based models in combination with advanced filtering techniques like the EnKF. For instance, it is a non-trivial question of how the observations relate to the agent states. In our case, we defined the H operator to be an aggregation operation, since agent states and observations both represented the same quantity, that is "wealth" measured in dollars. While this seems like a natural choice for micro-founded macro-economic models this might not always be possible. For instance, one might think about an agent-based model explaining inflation through labour market dynamics and central bank policy. Assume the inflation rate is the observation, yet the agents in the model, workers and banks perhaps, do not exhibit "inflation" as an individual characteristic and thus this cannot be simply be aggregated. Therefore, defining a more general framework of what observation operators for economic agent-based models look like is interesting for further studies. For comparison, work that has been conducted with the EnKF in urban system studies, and particularly pedestrian trajectory estimation, simply excludes variables that appear in the model but not in the observations through the observation operator (Suchak et al., 2024).

Another challenge with the wealth distribution particularly is that the right-tail of the distribution exhibits extreme values. This makes a smooth update experience difficult. The Kalman Gain weights for distinct wealth groups potentially develop in entirely different directions with very extreme absolute difference. Hence one natural variation of our work might be with log-transformed wealth values instead of untreated values. However for this the mathematics in Section 3.5 do not translate trivially. Again, our observation operator H is based on summing agent states to wealth groups and averaging them. Log-transformed values cannot be summed and averaged in the same way. A different operation other than basic matrix multiplication would be required for this to work.

In conclusion, our findings underscore the potential for economic models to harness data assimilation for more accurate and timely economic predictions, once challenges in data-model mapping are resolved.

CRediT authorship contribution statement

Yannick Oswald: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization. **Keiran Suchak:** Software, Validation, Writing – review & editing. **Nick Malleson:** Conceptualization, Validation, Resources, Writing – review & editing, Project administration, Funding acquisition.

Code and data accessibility

All code and data relating to this study is open-access at https://github.com/yannickoswald/real_time_ineq_abm.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 757455).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2024.106820.

Data availability

All code and data relating to this study is open-access at https://github.com/yannickoswald/real_time_ineq_abm.

References

Acemoglu, Daron, 2008. Introduction to Modern Economic Growth. Princeton University Press, p. 61.

Aktieninstitut, Deutsches, 2023. Shareholder numbers 2022: Germany can share! (in German)? URL: https://www.dai.de/en/detail/aktionaerszahlen-2022-trotz-rekordniveau-muss-die-politik-jetzt-liefern/.

Albert, Réka, Barabási, Albert-László, 2002. Statistical mechanics of complex networks. Rev. Modern Phys. 74 (1), 47.

Angle, John, 1986. The surplus theory of social stratification and the size distribution of personal wealth. Soc. Forces 65 (2), 293–326.

Athans, Michael, 1974. The importance of Kalman filtering methods for economic systems. In: Annals of Economic and Social Measurement, Volume 3, Number 1. NBER, pp. 49–64.

Auten, Gerald, Splinter, David, 2019. Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends. Working paper.

Axtell, Robert L., Farmer, J. Doyne, 2022. Agent-based modeling in economics and finance: Past, present, and future. J. Econ. Lit..

Blanchet, Thomas, Martinez-Toledano, Clara, 2023. Wealth inequality dynamics in Europe and the united states: Understanding the determinants. J. Monetary Econ. 133, 25–43.

Blanchet, Thomas, Saez, Emmanuel, Zucman, Gabriel, 2022. Real-Time Inequality. Tech. Rep., National Bureau of Economic Research.

Botta, Alberto, Caverzasi, Eugenio, Russo, Alberto, Gallegati, Mauro, Stiglitz, Joseph E., 2021. Inequality and finance in a rent economy. J. Econ. Behav. Organ. (ISSN: 0167-2681) 183, 998–1029. http://dx.doi.org/10.1016/j.jebo.2019.02.013, URL: https://www.sciencedirect.com/science/article/pii/S0167268119300447.

Caiani, Alessandro, Russo, Alberto, Gallegati, Mauro, 2019. Does inequality hamper innovation and growth? An AB-SFC analysis. J. Evol. Econ. 29 (1), 177–228. http://dx.doi.org/10.1007/s00191-018-0554-8, ISSN: 0936-9937, 1432-1386, (Visited 20 June 2024).

Chancel, Lucas, 2022. Global carbon inequality over 1990–2019. Nat. Sustain. 5 (11), 931–938.

Chancel, Lucas, Piketty, Thomas, Saez, Emmanuel, Zucman, Gabriel, 2022. World Inequality Report 2022. Harvard University Press.

Chiang, Yen-Sheng, 2015. Good samaritans in networks: An experiment on how networks influence egalitarian sharing and the evolution of inequality. PLoS One 10 (6), e0128777.

Clay, Robert, Ward, Jonathan A., Ternes, Patricia, Kieu, Le-Minh, Malleson, Nick, 2021. Real-time agent-based crowd simulation with the reversible jump unscented Kalman filter. Simul. Model. Pract. Theory (ISSN: 1569-190X) 113, 102386. http://dx.doi.org/10.1016/j.simpat.2021.102386, URL: https://www.sciencedirect.com/science/article/pii/S1569190X21000939.

Di Matteo, Tiziana, Aste, Tomaso, Hyde, S.T., 2003. Exchanges in complex networks: income and wealth distributions. arXiv preprint cond-mat/0310544.

Dragulescu, Adrian, Yakovenko, Victor M., 2000. Statistical mechanics of money. Eur. Phys. J. B 17, 723–729.

Drăgulescu, Adrian, Yakovenko, Victor M., 2001. Evidence for the exponential distribution of income in the USA. Eur. Phys. J. B 20, 585–589.

Evensen, Geir, 2003. The ensemble Kalman filter: Theoretical formulation and practical implementation. Ocean Dyn. 53, 343–367.

Evensen, Geir, 2009. The ensemble Kalman filter for combined state and parameter estimation. IEEE Control Syst. Mag. 29 (3), 83–104.

Fagereng, Andreas, Guiso, Luigi, Malacrino, Davide, Pistaferri, Luigi, 2020. Heterogeneity and persistence in returns to wealth. Econometrica 88 (1), 115–170.

Gilbert, G. Nigel, 2008. Agent-based models. Quantitative Applications in the Social Sciences, Sage Publications, Los Angeles, ISBN: 978-1-4129-4964-4, no. 07–153.

Hu, Xiaolin, 2022. Data assimilation for simulation-based real-time prediction/analysis. pp. 404–415.

Hu, Mao-Bin, Jiang, Rui, Wu, Yong-Hong, Wang, Ruili, Wu, Qing-Song, 2008. Properties of wealth distribution in multi-agent systems of a complex network. Phys. A 387 (23), 5862–5867.

Ichinomiya, Takashi, 2012. Wealth distribution on complex networks. Phys. Rev. E 86 (6), 066115.

Inglesi-Lotz, Roula, 2011. The evolution of price elasticity of electricity demand in South Africa: A Kalman filter application. Energy Policy 39 (6), 3690–3696.

Jones, Jeffrey M., 2023. What percentage of Americans own stock?. URL: https://news.gallup.com/poll/266807/percentage-americans-owns-stock.aspx.

Kalman, Rudolph Emil, 1960. A new approach to linear filtering and prediction problems.

Kalnay, Eugenia, 2003. Atmospheric Modeling, Data Assimilation and Predictability. Cambridge University Press.

Lee, Gyemin, Kim, Gwang Il, 2007. Degree and wealth distribution in a network induced by wealth. Phys. A 383 (2), 677–686.

Li, Hong, Kalnay, Eugenia, Miyoshi, Takemasa, 2009. Simultaneous estimation of covariance inflation and observation errors within an ensemble Kalman filter. Q. J. R. Meteorol. Soc. 135 (639), 523–533.

Lux, Thomas, 2005. Emergent statistical wealth distributions in simple monetary exchange models: a critical review. In: Econophysics of Wealth Distributions: Econophys-Kolkata I. Springer, pp. 51–60.

Lux, Thomas, 2018. Estimation of agent-based models using sequential Monte Carlo methods. J. Econom. Dynam. Control 91, 391–408.

Malleson, N., Minors, K., Kieu, L.-M., Ward, J.A., West, A., Heppenstall, A., 2020. Simulating crowds in real time with agent-based modelling and a particle filter. J. Artif. Soc. Soc. Simul. 23 (3), 3. http://dx.doi.org/10.18564/jasss.4266.

Mattison, Siobhán M, Smith, Eric A, Shenk, Mary K, Cochrane, Ethan E, 2016. The evolution of inequality. Evol. Anthropol. 25 (4), 184–199.

Monti, Corrado, Pangallo, Marco, De Francisci Morales, Gianmarco, Bonchi, Francesco, 2023. On learning agent-based models from data. Sci. Rep. 13 (1), 9268.

Munguia, Rodrigo, Davalos, Jessica, Urzua, Sarquis, 2019. Estimation of the Solow-Cobb-Douglas economic growth model with a Kalman filter: An observability-based approach. Heliyon 5 (6).

Palagi, Elisa, Napoletano, Mauro, Roventini, Andrea, Gaffard, Jean-Luc, 2023. An agent-based model of trickle-up growth and income inequality. Econ. Model. 129, 106535.

Papadopoulos, Georgios, 2019. Income inequality, consumption, credit and credit risk in a data-driven agent-based model. J. Econom. Dynam. Control 104, 39–73.

Pasricha, Gurnain Kaur, 2006. Kalman filter and its economic applications. In: Munich Personal RePEc Archive.

Poledna, Sebastian, Miess, Michael Gregor, Hommes, Cars, Rabitsch, Katrin, 2023. Economic forecasting with an agent-based model. Eur. Econ. Rev. 151, 104306.

Schneider, Wolfgang, 1988. Analytical uses of Kalman filtering in econometrics—A survey. Statist. Papers 29 (1), 3–33.

Semieniuk, Gregor, 2024. Inconsistent definitions of GDP: Implications for estimates of decoupling. Ecol. Econom. 215, 108000.

Silva, Petrônio CL, Batista, Paulo VC, Lima, Hélder S, Alves, Marcos A, Guimarães, Frederico G, Silva, Rodrigo CP, 2020. COVID-abs: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions. Chaos Solitons Fractals 139, 110088.

Souma, Wataru, Fujiwara, Yoshi, Aoyama, Hideaki, 2003. Wealth distribution in scale-free networks. In: Meeting the Challenge of Social Problems Via Agent-Based Simulation: Post-Proceedings of the Second International Workshop on Agent-Based Approaches in Economic and Social Complex Systems. Springer, pp. 37–49.

Suchak, Keiran, Kieu, Minh, Oswald, Yannick, Ward, Jonathan A, Malleson, Nick, 2024. Coupling an agent-based model and an ensemble Kalman filter for real-time crowd modelling. Royal Soc. Open Sci. 11 (4), 231553.

Swarup, Samarth, Mortveit, Henning S., 2020. Live simulations. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems. In: AAMAS '20, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, ISBN: 978-1-4503-7518-4, pp. 1721–1725.

Talagrand, O., 1991. The use of adjoint equations in numerical modelling of the atmospheric circulation. In: Griewank, A., Corliss, G.F. (Eds.), Automatic Differentiation of Algorithms: Theory, Implementation, and Application. SIAM, Philadelphia, PA, ISBN: 0-89871-284-X, pp. 169–180.

Tang, Daniel, Malleson, Nick, 2022. Data assimilation with agent-based models using Markov chain sampling. arXiv preprint arXiv:2205.01616.

Ternes, Patricia, Ward, Jonathan A, Heppenstall, Alison, Kumar, Vijay, Kieu, Le-Minh, Malleson, Nick, 2022. Data assimilation and agent-based modelling: towards the incorporation of categorical agent parameters. Open Res. Eur. 1.

Thamae, Retselisitsoe Isaiah, Thamae, Leboli Zachia, Thamae, Thimothy Molefi, 2015. Dynamics of electricity demand in lesotho: A Kalman filter approach. Stud. Bus. Econ. 10 (1), 130–139.

Vallejos, Hunter A., Nutaro, James J., Perumalla, Kalyan S., 2018. An agent-based model of the observed distribution of wealth in the United States. J. Econ. Interact. Coord. 13, 641–656.

Vishwakarma, Keshav P., 1970. Prediction of economic time-series by means of the Kalman filter. Int. J. Syst. Sci. 1 (1), 25–32.

Wang, M., Hu, X., 2015. Data assimilation in agent based simulation of smart environments using particle filters. Simul. Model. Pract. Theory 56, 36–54. http://dx.doi.org/10.1016/j.simpat.2015.05.001.

Ward, Jonathan A., Evans, Andrew J., Malleson, Nicolas S., 2016. Dynamic calibration of agent-based models using data assimilation. Royal Soc. Open Sci. 3 (4), 150703.

World Inequality Database, 2024. World inequality database. https://wid.world. (Accessed 25 January 2024).

Yakovenko, Victor M., Rosser, Jr., J. Barkley, 2009. Colloquium: Statistical mechanics of money, wealth, and income. Rev. Modern Phys. 81 (4), 1703.

Yang, Xiaoliang, Zhou, Peng, 2022. Wealth inequality and social mobility: A simulation-based modelling approach. J. Econ. Behav. Organ. 196, 307–329.

Zhou, Tengjiao, Ding, Long, Ji, Jie, Li, Lin, Huang, Weiwei, 2019. Ensemble transform Kalman filter (ETKF) for large-scale wildland fire spread simulation using FARSITE tool and state estimation method. Fire Saf. J. 105, 95–106.