NBER WORKING PAPER SERIES

SELF-FULFILLING PROPHECIES, QUASI NON-ERGODICITY & WEALTH INEQUALITY

Roger Farmer
Jean-Philippe Bouchaud

Working Paper 28261
http://www.nber.org/papers/w28261

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2020

This paper was written after J. Doyne Farmer suggested that we collaborate as co-leaders of the Instability Hub for the ESRC funded Network Plus, Rebuilding Macroeconomics. We thank Angus Armstrong, Pablo Beker, Michael Benzaquen, Leland E. Farmer, Alan Kirman, Robert McKay, Ian Melbourne, José Moran, Patrick Pintus and Ole Peters for many insightful discussions on the topic of the paper. Thanks to C. Roxanne Farmer for helpful suggestions. We reserve a special thanks to J. Doyne Farmer for his insight that two people with such disparate backgrounds would have something to learn from each other. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Self-Fulfilling Prophecies, Quasi Non-Ergodicity & Wealth Inequality
Roger Farmer and Jean-Philippe Bouchaud
NBER Working Paper No. 28261
December 2020
JEL No. E0,G0,G12,G14

ABSTRACT

We construct a model where people trade assets contingent on an observable signal that reflects public opinion. The agents in our model are replaced occasionally and each person updates beliefs in response to observed outcomes. We show that the distribution of the observed signal is described by a quasi non-ergodic process and that people continue to disagree with each other forever. Our model generates large wealth inequalities that arise from the multiplicative nature of wealth dynamics which makes successful bold bets highly profitable. The flip side of this statement is that unsuccessful bold bets are ruinous and lead the person who makes such bets into poverty. People who agree with the market belief have a low expected subjective gain from trading. People who disagree may either become spectacularly rich, or spectacularly poor.

Roger Farmer
University of California, Los Angeles
Department of Economics
Box 951477
Los Angeles, CA 90095-1477
and CEPR
and also NBER
rfarmer@econ.ucla.edu

Jean-Philippe Bouchaud
23 Rue de L'Université
75007 Paris
Jean-Philippe.BOUCHAUD@cfm.fr

Self-Fulfilling Prophecies, Quasi Non-Ergodicity & Wealth Inequality

By Jean-Philippe Bouchaud and Roger E.A. Farmer*

We construct a model where people trade assets contingent on an observable signal that reflects public opinion. The agents in our model are replaced occasionally and each person updates beliefs in response to observed outcomes. We show that the distribution of the observed signal is described by a quasi non-ergodic process and that people continue to disagree with each other forever. Our model generates large wealth inequalities that arise from the multiplicative nature of wealth dynamics which makes successful bold bets highly profitable. The flip side of this statement is that unsuccessful bold bets are ruinous and lead the person who makes such bets into poverty. People who agree with the market belief have a low expected subjective gain from trading. People who disagree may either become spectacularly rich, or spectacularly poor.

I. Introduction

A. Quasi Non-Ergodic Economies

Almost all quantitative work in macroeconomics and finance assumes that economic and financial time series can be represented by an ergodic stochastic process. Ergodicity means that the average over a single long trajectory is equal to the average over multiple realisations of trajectories drawn from the same process. In this paper we construct a model of asset prices where this assumption holds over very long time periods, but breaks down for time periods that might realistically characterize the decision window of any living human being. Processes of this kind are referred to in the natural sciences as *quasi-non-ergodic*.[1]

Our work is closely related to the Pólya urn model reviewed in Pemantle (2007). In this model, an urn contains M red balls and $(N-M)$ black balls. A ball is

* Bouchaud: Capital Fund Management and Institut de France Jean-Philippe.Bouchaud@academie-sciences.fr. Farmer: Department of Economics, University of Warwick and Department of Economics, UCLA, r.farmer.1@warwick.ac.uk. This paper was written after J. Doyne Farmer suggested that we collaborate as co-leaders of the Instability Hub for the ESRC funded Network Plus, Rebuilding Macroeconomics. We thank Angus Armstrong, Pablo Beker, Michael Benzaquen, Leland E. Farmer, Alan Kirman, Robert McKay, Ian Melbourne, José Moran, Patrick Pintus and Ole Peters for many insightful discussions on the topic of the paper. Thanks to C. Roxanne Farmer for helpful suggestions. We reserve a special thanks to J. Doyne Farmer for his insight that two people with such disparate backgrounds would have something to learn from each other.

[1] Quasi-non-ergodicity is common to many physical systems where ergodicity appears to be broken at intermediate times but is restored at very long (sometimes astronomical) times, as in the case of glasses and "spin-glasses" (Debenedetti and Stillinger, 2001; Anderson, 1989).

chosen at random. If it is red (respectively black) then 2 red balls (respectively black balls) are re-introduced in the urn, which now contains $N+1$ balls. The probability of drawing a ball with a specific colour therefore increases with the number of times this colour was selected in the past. The long term fate of this simple dynamic is surprising. As $t \to \infty$, the probability $\mathbb{P}_t$ to draw a red ball converges to some limiting value $\mathbb{P}_\infty$, but the value of this asymptotic probability is itself random. Starting from the very same urn with $N_0/2$ red balls and $N_0/2$ black balls, two different runs of the dynamics will lead to two different values of $\mathbb{P}_\infty$, although each of them becomes stable over time, corresponding to a stationary process. To completely characterize the behaviour of this process one must introduce probabilities over probabilities.[2] The Pólya urn model is an example of a non-ergodic stochastic process.

We construct an economic model where a mechanism similar to the one at play in Pólya urns is present. In our model, economic agents attempt to predict a binary outcome $s \in \mathbf{S} \equiv \{0,1\}$ which might represent, for example, high or low output in the next month. Each agent has a time dependent belief of the probability of the good outcome which she updates from observation. Importantly, we assume that the true probability of the realized outcome is equal to the average of the subjective probabilities of the agents.

We show that the coupling of individual and aggregate beliefs leads to a process in which the true probability of a high outcome wanders randomly in the unit interval. The conditional distribution of this random variable converges, in the limit, to an invariant unconditional probability distribution which is uniform on $[0,1]$. Although the unconditional probability of a good outcome is always one half, it is always better to use recent information to predict the next draw than to use very long time-series averages. Our self-referential learning model is an example of a quasi-non-ergodic stochastic process.

B. Beliefs and Wealth Inequality

In many economic models where people begin with different prior beliefs, their belief differences do not survive repeated interactions in markets. That turns out, in our model, not to be true. We endow our probabilistic world with a market that allows agents to trade two securities. One security pays one unit of consumption if and only if $s = \{1\}$ and a second security pay one unit if and only is $s = \{0\}$. Securities with these characteristics are referred to as *Arrow securities* (Arrow, 1964). Agents solve an inter-temporal optimisation programme to determine how much of each security they wish to hold at time t and, because agents have different beliefs, they are willing to trade.

Naively, one would expect the market clearing price should reveal the average

[2]Whereas the precise value of $\mathbb{P}_\infty$ is unpredictable, the probability $\mathcal{P}$ of the probability $\mathbb{P}_\infty$ is a $\beta(a,b)$ distribution over $[0,1]$ where the parameters a and b depend on the initial distribution of balls. Starting from an urn with two balls, one red and one black, $\mathbb{P}_\infty$ has a $\beta(1,1)$ distribution which is uniform on $[0,1]$ implying that the limiting probability, $\mathbb{P}_\infty$, can be anywhere in $[0,1]$ with equal probability.

belief and cause traders to coordinate on the true probability. But in our model this is not the case. We show that the probability implied by market prices is a *wealth-weighted* average of individual subjective beliefs. The wealth of each agent fluctuates with time and some people accidentally, and temporarily, predict the true probability better than others. Interestingly, the distribution of wealth is so unequal that the probability implied by market prices is dominated by the wealthiest agents and fails to reveal the true probability of the good outcome. Because markets fail to aggregate private information correctly, they cannot be used by individuals to reveal the truth.

C. Literature Review

There is an extensive literature on self-fulfilling prophecies in rational expectations models. Early versions of this literature that rely on dynamic indeterminacy are discussed in Farmer's (1999) textbook and more recent models that display hysteresis and steady-state indeterminacy are reviewed in Farmer (2020).[3] The literature on self-fulfilling prophecies explains how beliefs drive economic fluctuations, but as with all rational expectations models, eventually everybody agrees with everybody else. Our current paper, in contrast, explains how a large number of agents interacting in a complete set of financial markets can continue to disagree forever.

The title of our paper, which features the concept of quasi non-ergodicity, is inspired by the observation that although ergodicity may be a feature of very long sequences of random variables, it may not hold on time scales relevant to the lifetimes of economic decision makers. Our model is a close cousin of Kirman's ant model (Kirman, 1993), also known as the Moran model in the theory of population dynamics (Moran, 1958).

We are not the first to explore the topic of non-ergodicity for economics. Brock and Durlauf (2001) have shown that interaction effects can trap the economy in a path-dependent state. Horst (2017) reviews the literature on path dependence and, more recently, Moran et al. (2020b) show that ergodicity breaking occurs in models of habit formation. In contrast to these papers, we focus on a case where ergodicity is not strictly broken but where the time scale over which it applies may be astronomically long.

Peters (2019) has pointed out that identifying time averages over a single trajectory with ensemble averages can lead to misleading conclusions, and that special care should be devoted to the choice of an appropriate, process dependent, utility function. Our model, on the other hand, illustrates a rather different facet of non-ergodicity, where agents adapt their beliefs based on a observation window much shorter than the time needed to reach ergodicity.

Our paper deals with the asymptotic properties of a multi-agent complete-market economy where agents have heterogeneous priors. Previous related work

[3]The importance of correctly modelling beliefs in models of self-fulfilling prophecies is further discussed in Farmer (2021).

includes Blume and Easley (2006) who study the asymptotic properties of consumption and Cogley and Sargent (2008, 2009) who study asset prices in an economy with informed and uninformed agents. The closest previous paper to ours is by Beker and Espino (2011) who study a stochastic endowment economy populated by infinitely-lived Bayesian updaters.

In Beker and Espino (2011), agents have different priors but, in their central case everyone eventually learns the truth. We modify the Beker-Espino environment in two ways. First, the process that generates the states is self-referential and leads to a quasi non-ergodic process. Second, we modify the environment to allow replacement of agents and we endow new agents with a random prior. Our work is similar to the stochastic extensions by Farmer et al. (2011) and Farmer (2018) of Blanchard's (1985) perpetual youth model. The replacement of agents with new people with random priors is central to our demonstration that beliefs never converge.

A second important assumption that drives our results is the assumption that agents use constant gain learning to update their beliefs. Adam et al. (2016) and Adam et al. (2017) also drop Bayesian updating and use constant gain learning. Unlike those papers, we study a multi-agent economy and we link the true stochastic process to subjective beliefs through a beauty contest game. As a consequence, the event probability itself is time dependent in our model, and the wealth distribution is non-trivial, even in the limit.

D. Fat-tailed Wealth Distributions and Wealth Inequality

In most countries, the wealth distribution has *fat tails*, a.k.a. Pareto tails Gabaix (2009); Piketty and Zucman (2014). One measure of this property is described by a log-log plot of the complementary cdf of wealth on the y-axis against wealth on the x-axis, for all values of wealth greater than the mode. The slope of this plot provides an estimate of the rate at which the right tail of the wealth distribution decays with wealth. If the absolute value of this slope is between 0 and 1, there is a well defined invariant wealth distribution, but the mean of this distribution does not exist. If the absolute value of this slope is between 1 and 2, the mean of the wealth distribution is finite, but the second and all higher moments do not exist. Empirical estimates of the tail parameter of the U.S. wealth distribution are centered around 1.5. This fact implies that there exists a well defined wealth distribution with a finite mean, but the variance and higher moments do not exist.

One might think that the wealth distribution is unequal because people are endowed with different abilities and some individuals are able to exploit these abilities to earn more than others. Although the income distribution, like the wealth distribution, has a fat right tail, the magnitude of the tail parameter is too large to explain wealth inequality.[4] This fact suggests that we should look

[4]Benhabib and Bisin (2011) show that some simple economic theories constrain the tail parameter on

elsewhere for an explanation of wealth inequality.

In their survey of this topic, Benhabib and Bisin (2011) point out that models where people earn different rates of return on the same asset can explain a fat tailed wealth distribution (see also Bouchaud and Mézard (2000); Gabaix et al. (2016)). Our model provide a mechanism where this happens in equilibrium as a consequence of persistent differences in individual assessments of risk. In our model, everyone has the same income but the wealth distribution is highly unequal as a consequence of the multiplicative nature of wealth accumulation.

II. A Two-Outcome, Self-Referential Model

We will build up our argument in three stages. In stage one, we describe a game in which agents form beliefs about a binary outcome and we show that our game leads to a quasi-non-ergodic process for the true belief. In stage two, we embed our agents in an endowment economy and we allow them to trade Arrow securities contingent on the realization of the binary random variable. In stage three, we show that the contingent securities market can be replaced by debt and equity and that the equilibria of this more realistic version of our model is the same as the model in which agents trade Arrow securities.

A. *The Beauty Contest Game*

We assume that N agents play a game in which each person must forecast the average belief of the other agents about the outcome of a sequence of binary random events $\{s_t \in \mathbf{S}\}_{t=1}^{\infty}$. This is a simple version of a game that Keynes famously used in *The General Theory* (Keynes, 1936) to motivate his view that the stock market is driven by what he called 'animal spirits'.

We represent the belief held at date $t-1$ of agent i of the probability that $s_t = \{1\}$ as $\mathbb{P}_{i,t}(s = \{1\})$ and we model the self-referential nature of beliefs by assuming that the *true probability* of the event, $\mathbb{P}_t(s = \{1\})$, is equal to the average belief,[5]

$$\mathbb{P}_t \equiv \sum_{i=1}^{N} \frac{\mathbb{P}_{i,t}}{N}. \tag{1}$$

where throughout the paper, we will drop the argument $s = \{1\}$ after $\mathbb{P}$, unless we explicitly need to distinguish the two outcomes.

People communicate with each other on social networks and a central player generates the random variable s_t which takes the value $\{1\}$ with probability $\mathbb{P}_t$

the wealth distribution to be *greater* than the tail parameter on the income distribution, implying that wealth should be *less* unequally distributed than income. But estimates by e.g. Badel et. al. (2017) of the tail parameter for the U.S. income distribution are close to 2; a number greater than 1.5. Similarly the Gini coefficient for income is always found to be smaller than the one for wealth.

[5]More generally, one can consider a model where the true probability is a non linear, sigmoidal function $\mathcal{F}$ of the average belief: see Appendix A.A2.

and $\{0\}$ with probability $1 - \mathbb{P}_t$. At this point this is an abstract game with no economic consequences.

One interpretation of our model is that $\mathbb{P}_{i,t}$ is the degree of confidence that people have in the future. A mechanism by which this confidence impacts real outcomes may be the following. People communicate their beliefs through social interaction and the central player in our game represents an influential financial journalist. The journalist aggregates information and, with probability $\mathbb{P}_t$ he writes an article with a positive outcome and, with probability $1 - \mathbb{P}_t$ he writes an article with a negative outcome.

B. *How People Update Beliefs*

The properties of an economic model will depend heavily on the assumptions we make about how the players change over time. One can show that, if agents are infinitely lived Bayesian updaters with different priors, that $\mathbb{P}$ converges to a number in $[0, 1]$, but that number is different for every realization of $\{\mathbb{P}_t\}_{t=1}^{\infty}$. This setup is an economic analog of the Pólya urn model and although the example is instructive, it is not very interesting as a theory of asset market trade. Everyone eventually converges on the truth and although the truth is itself a function of history, eventually people all agree with one another.

To generate a theory of permanent disagreement we modify the model in two ways. First, we replace the assumption of Bayesian updating with an alternative constant gain learning algorithm in which people discount the past at rate λ. Second, we allow the set of decision makers to change over time by recognizing that people have finite lives. We assume that people die with probability δ but the death probability is independent of age. When a person dies, she is replaced by a new person with belief $\mathbb{P}_i = z_i$ where z_i is a random variable drawn from a uniform measure on $[0, 1]$. In Section III we will use these assumptions to generate simple expressions for aggregate asset prices in a market economy.[6]

We keep track of who lives and who dies by introducing a random vector $\boldsymbol{x}_t \in \mathbf{X} \equiv \{0, 1\}^N$, where $x_{i,t} = 1$ with probability $1 - \delta$ and 0 with probability δ. If a person who was alive in period $t - 1$ survives into period t then $x_{i,t} = 1$. If she dies then $x_{i,t} = 0$. We assume that the evolution of the beliefs of the person with index i is given by the expression

$$\mathbb{P}_{i,t+1} = x_{i,t}[(1 - \lambda)\mathbb{P}_{i,t} + \lambda s_t] + (1 - x_{i,t})z_{i,t}, \tag{2}$$

where $\boldsymbol{z}_t \in \mathbf{Z} \equiv [0, 1]^N$ and each of element of $\boldsymbol{z}_t$ is an independent draw from a uniform distribution.[7]

[6]Our model is a version of the perpetual youth model of Blanchard (1985) as extended to the stochastic case by Farmer et al. (2011) and Farmer (2018).

[7]The exact form of the distribution of z_t is not important for any of our results. One could also assume that $z_{i,t}$ is a weighted sum of the average belief $\mathbb{P}_t$ and a uniform random variable. This would not change the structure of the model at all, only the meaning of the parameters.

The term in square brackets on the right side of Eq. (2) represents the way that a person who is alive in two consecutive periods updates her belief. She uses constant gain learning with gain parameter λ where a value of λ closer to 1 means that the person puts more weight on recent outcomes. This term is multiplied by $x_{i,t}$ to reflect the fact that it applies only if person i survives into the period. The second term on the right side of Eq. (2) is multiplied by $1 - x_{i,t}$. This reflects the assumption that if agent i dies, her position is filled by a new-born person who starts life with a random subjective belief, $z_{i,t}$. We assume that $z_{i,t}$ is distributed uniformly on $[0, 1]$.

C. The Behaviour of Beliefs in the Large N Limit

In Section V we will use Eq. (2) to simulate data from an economy with a large but finite number of agents. To better understand the properties of those simulations, in this subsection we study the properties of Eq. (2) as $N \to \infty$ and as the length of a period converges to zero. We refer to this case as the continuous-time, large N limit. First, we retain the discrete time assumption, and study the behaviour of Eq. (2) for large N. We refer to this as the discrete-time, large N limit.

To arrive at expressions for the discrete-time large N limit we combine equations (1) and (2) and we take $N \to \infty$. This leads to the expression,

$$\mathbb{P}_{t+1} = (1-\delta)\left[(1-\lambda)\mathbb{P}_t + \lambda s_t\right] + \delta \mathbb{E}[z], \tag{3}$$

where $\mathbb{E}[z] = \frac{1}{2}$. Eq. (3) defines a Markov process for the random variable $\mathbb{P}$ with a transition operator $\mathcal{T}[\mathcal{P}]$ for its probability $\mathcal{P}$ that is defined by the integral equation,

$$\begin{aligned} \mathcal{T}[\mathbb{P}] \equiv \int_0^1 \mathrm{d}\mathbb{P}\, \mathcal{P}(\mathbb{P}) \Bigg[\mathbb{P}\, \mathbf{d}\left(\mathbb{P}' - (1-\delta)[(1-\lambda)\mathbb{P} + \lambda] - \frac{\delta}{2}\right) \\ + (1-\mathbb{P})\, \mathbf{d}\left(\mathbb{P}' - (1-\delta)(1-\lambda)\mathbb{P} - \frac{\delta}{2}\right)\Bigg], \end{aligned} \tag{4}$$

where $\mathbf{d}(\cdot)$ is Dirac's delta function.

Let $\mathcal{P}_t(\mathbb{P})$ be the probability density at date t that $\mathbb{P}_{t+1}$ takes any given value in $[0, 1]$. Then

$$\mathcal{P}_{t+1}(\mathbb{P}) = \mathbb{T}\mathcal{P}_t(\mathbb{P}),$$

is the probability density at date $t + 1$ and

$$\mathcal{P}_T(\mathbb{P}) = \mathcal{T}^T \mathcal{P}_0(\mathcal{P}),$$

is the probability density at date T where $\mathcal{T}^T$ is the T'th iterate of the operator $\mathcal{T}$. Notice that $\mathcal{P}$ defines a probability density over probabilities.

Introducing the change of variable u such that $u = \mathbb{P} - \frac{1}{2}$, we show in Appendix A that in the continuous time limit, $\mathcal{P}_t(u)$ converges to a symmetric beta-distribution with parameter $\alpha = \delta/\lambda^2$,

$$\mathcal{P}_\infty(u) = \frac{\Gamma(2\alpha)}{\Gamma^2(\alpha)} \left(\frac{1}{4} - u^2\right)^{\alpha-1}. \tag{5}$$

This distribution is U-shaped when $\delta < \lambda^2$ and hump-shaped for $\delta > \lambda^2$. In our simulations we calibrate the model by choosing $\delta = \lambda^2$, a case for which the invariant measure is uniform on $[0, 1]$.[8]

In the continuous time limit $(\lambda, \delta \to 0)$, we are able to derive an exact expression for the degree of disagreement which we define as the difference between the belief of agent i and the average belief. In symbols,

$$\mathbb{D}_t \equiv \mathbb{P}_{i,t} - \mathbb{P}_t.$$

Combining this definition with equations (2) and (3) and rearranging terms leads to the following equation which determines the stochastic evolution of $\mathbb{D}_{i,t+1}$,

$$\begin{aligned}\mathbb{D}_{i,t+1} = x_{i,t} &\left[(1-\lambda)\mathbb{D}_{i,t} + \delta\left((1-\lambda)\mathbb{P}_t + \lambda(s_t - \frac{1}{2})\right)\right] \\ &+ (1 - x_{i,t})\left[z_i - (1-\delta)\left((1-\lambda)\mathbb{P}_t + \lambda s_t + \frac{\delta}{2}\right)\right].\end{aligned} \tag{6}$$

We show in Appendix B, that the unconditional expectation of $\mathbb{D}_i$, converges to zero almost surely and that, in the large N, small λ limit its variance is given by the expression,

$$\mathbb{V}[\mathbb{D}_i] = \left[\frac{\lambda}{2 + (1-\alpha)\lambda}\right] \frac{\alpha(\alpha+2)}{6(2\alpha+1)} + O(\lambda^3), \qquad \alpha = \frac{\delta}{\lambda^2}. \tag{7}$$

The standard deviation of $\mathbb{D}_i$ is a measure of disagreement between agents in the unconditional limiting distribution. For our calibration, $\alpha = 1$ and for this parameterisation

$$\mathbb{V}[\mathbb{D}_i] \approx \frac{\sqrt{\delta}}{12}, \qquad (\delta \ll 1). \tag{8}$$

This expression implies that the typical disagreement between agents, as measured by the standard deviation of $\mathbb{D}_i$, only vanishes at rate $\delta^{1/4}$, as $\delta \to 0$, i.e. very slowly. Even for long lifetimes, people still disagree significantly.

[8] When $\delta \to 0$, the distribution of $\mathbb{P}$ becomes highly peaked around 0 and 1. This model is studied in detail in Moran et al. (2020a) who show that the time spent by $\mathbb{P}_t$ in the vicinity of 0 or 1 is equal to δ^{-1} and is independent of λ. This *switching time* also corresponds to the *ergodic time* defined as the time required for $\mathcal{P}$ to approach the stationary distribution $\mathcal{P}_0$.

In our simulations we chose a time interval of one week and we set $\delta = 3.9\times10^{-4}$. These choices imply that life expectancy, averaged over people of all ages, is approximately 50 years which accords well with crude estimates from US actuarial tables. For this calibration, the standard deviation of $\mathbb{D}_i$ is approximately 4% for $\alpha = 1$. This level of disagreement between agents is quite reasonable and, as we will show in our simulations, large enough to generate substantial discrepancies between the market price and the true price, and a "fat" power-law right tail of the wealth distribution when people make bets based on their subjective beliefs.

III. Heterogeneous Beliefs in a Market Economy

We have built a model to describe the evolution of aggregate opinions. But what happens if people trade with other people with different beliefs? To answer that question we construct an endowment economy where each person is endowed with ε units of a non-storable commodity in every period in which she is alive. We further assume that people trade a complete set of Arrow securities, indexed to the *exogenous state*, which we represent by σ. We use the adjective *exogenous*, to distinguish the vector σ from a vector of *endogenous states* that we introduce in Section III.C.

A. *The Definition of the Exogenous State*

The exogenous state has three elements. The first element, $s \in \mathbf{S} \equiv \{0,1\}$, is the realization of a public signal. The second element, $\boldsymbol{x} \in \mathbf{X} \equiv \{0,1\}^N$, is a vector that differentiates newborns from survivors and the third element, $\boldsymbol{z} \in \mathbf{Z} \equiv [0,1]^N$, encodes the conditional probabilities of newborns.[9] Putting these pieces together we have that $\sigma \equiv \{s, \boldsymbol{x}, \boldsymbol{z}\} \in \boldsymbol{\Sigma} \equiv \mathbf{S} \times \mathbf{X} \times \mathbf{Z}$. We use a prime to denote the state in period $t+1$.

At each date, people trade a complete set of Arrow securities which depend not just on s', but also on the realizations of $\boldsymbol{x}'$ which encodes who lives and who dies. There are 2^N possible realizations of $\boldsymbol{x}'$ where the i'th element of $\boldsymbol{x}'$ equals $\{1\}$ if person i survives and $\{0\}$ if she dies. The $\sigma' = (s', \boldsymbol{x}')$ security costs $Q(\sigma'|\sigma)$ commodities at date t and pays 1 commodity at date $t+1$ if and only if state σ' occurs. We assume that everybody knows the probability of birth and death of everyone alive today but they have different beliefs, represented by $\mathbb{P}_i$, of the probability that $s' = \{1\}$.

We refer to a realization $\boldsymbol{x}'$ as a *mortality state* and we denote the probability of a realization of $\boldsymbol{x}'$ by $p(\boldsymbol{x}')$. We assume that $p(\boldsymbol{x}')$ is common knowledge and that $\boldsymbol{x}'$ is independent of s'. These assumptions allow us to factor $\mathbb{P}_i(\sigma')$ into two components, $\mathbb{P}_i(s')$, which is person i's subjective conditional probability that

[9] We generate this vector for all i, including survivors from the previous period. Notice, however, that z_i only enters the model when multiplied by $1 - x_i$ which is zero for survivors.

$s' = \{1\}$ and $p(\boldsymbol{x}')$, which is the objective probability of the mortality state.

$$\mathbb{P}_i(\sigma') = \mathbb{P}_i(s')p(\boldsymbol{x}'). \tag{9}$$

This completes our definition of the exogenous state. In the subsequent subsection we define the objectives and constraints of individual agents and we derive a set of rules that represents their behaviour in an exchange economy.

B. A Model of Rational Choice

We assume that agents maximize the discounted expected utility of the logarithm of consumption. This assumption implies that our agents choose to spend a fixed fraction of wealth in each period on the consumption good. The novel aspect of our approach is the decision rule we derive which shows how agents allocate their wealth to the two Arrow securities. This decision rule depends on their subjective beliefs, which evolve in the manner described in Section II.B.

First, we break wealth into two components; human wealth and financial wealth. The *human wealth* of person i is defined by the recursion,

$$H_i(\sigma) = \varepsilon + \sum_{\sigma'} Q(\sigma'|\sigma)\, x_i'\, H_i(\sigma'), \tag{10}$$

and *aggregate human wealth* $H(\sigma)$ is the sum of individual human wealth over all living persons,

$$H(\sigma) = N(\sigma)H_i(\sigma). \tag{11}$$

Next, we define *financial wealth* of person i, $a_i(\sigma)$, to be the value of Arrow securities brought into period t. The *total wealth* of person i is the sum of human wealth and *financial wealth*

$$W_i(\sigma) = H_i(\sigma) + a_i(\sigma). \tag{12}$$

Each period, the agent faces the following budget equation,

$$\sum_{\sigma'} x_i'(\sigma')Q(\sigma'|\sigma)a_i'(\sigma') + c_i(\sigma) = a_i(\sigma) + \varepsilon. \tag{13}$$

The right side of Eq. (13) represents a person's available resources at date t. The left side represents the ways those resource can be allocated; to consumption or to the accumulation of a bundle of Arrow securities that will be available for consumption or saving in the subsequent period.

We model the consumption and asset allocations of each person as the unique solution to the following maximization problem:

PROBLEM 1:

$$V_i[W_i(\sigma)] = \max_{W_i'(\sigma')} \left[\log c_i(\sigma) + \beta \sum_{\sigma'} \mathbb{P}_i(\sigma') x_i'(\sigma') V_i'[W_i'(\sigma')] \right] \tag{14}$$

such that

$$\mathbb{P}_i(\sigma') = x_i'[(1-\lambda)\mathbb{P}_i(\sigma) + \lambda s] + (1 - x_i') z_i', \tag{15}$$

and

$$\sum_{\sigma'} x_i(\sigma') Q(\sigma'|\sigma) W_i(\sigma') + c_i(\sigma) \le W_i(\sigma). \tag{16}$$

In Section II.B we derived an expression for the evolution of person i's beliefs. Eq. (15) reproduces that equation using the definition of σ and replacing time subscripts with prime notation.

Eq. (16) is derived by combining equations (10) and (13). $V_i[W_i(\sigma)]$ is the maximum attainable utility given wealth $W_i(\sigma)$, $c_i(\sigma)$ is date t consumption and β is the common discount rate. Following common usage we refer to the consumption decision that solves Problem 1 as the *policy function* and to the maximum attainable utility as a function of wealth as the *value function*.

PROPOSITION 1: *The policy function and the value function for Problem 1 are given by Equations* (17) *and* (18)*,*

$$c_i(\sigma) = [1 - \beta(1-\delta)] W_i(\sigma), \tag{17}$$

$$V_i[W_i(\sigma)] = \frac{1}{1 - \beta(1-\delta)} \log[W_i(\sigma)] + B, \tag{18}$$

where B is a constant that can be computed but its value is irrelevant for our purpose.

The wealth of the person with label i evolves according to Eq. (19)

$$W_i(\sigma') = x_i' \left[\frac{\beta \mathbb{P}_i(\sigma')}{Q(\sigma'|\sigma)} W_i(\sigma) \right] + (1 - x_i') H_i(\sigma'), \tag{19}$$

where $H(\sigma)$ is defined by the recursion Eq. (10)*.*

The first term on the right side of Eq. (19) is the wealth evolution equation for person i if she survives into period $t+1$. The second term on the right side of the equation resets person i's wealth to $H(\sigma')$ if she dies and is replaced by a newborn. For a proof of Proposition 1, see Appendix C.

C. Definition of Equilibrium

We have constructed a theory of individual choice. According to this theory, peoples' decisions are a function of the exogenous state and of the stochastic process for prices. In this section we construct an equilibrium theory where prices are determined to set the excess demands for goods and the excess demands for Arrow securities, in every period, to zero. First, we define a new object; the *endogenous state.*

The endogenous state has two elements. The first element, $P \in \mathbf{P} \equiv [0,1]^N$ is a vector of subjective conditional probabilities with generic element $\mathbb{P}_i$. The second element, $W \in \mathbf{W} \equiv \mathbb{R}_+^N$ is a vector of wealth positions with generic element W_i. Putting these pieces together, the endogenous state is represented by $y \equiv \{P, W\} \in \mathbf{Y} \equiv \mathbf{P} \times \mathbf{W}$.

Next, we derive a function $\mathcal{G}(\cdot)$ to explain how the endogenous state evolves through time. Our approach is a relatively standard application of recursive equilibrium theory (Stokey et al., 1989). Our innovation, over conventional dynamic stochastic general equilibrium models, is to provide a self-referential theory of learning in which the economy does not converge to a rational expectations equilibrium.

We begin with a definition of recursive equilibrium:

DEFINITION 1 (Recursive Equilibrium): *A recursive equilibrium is a price function $Q : \mathbf{\Sigma}^2 \to \mathbf{Q} \equiv [0,1]^{2N}$ and a state evolution function $\mathcal{G} : \mathbf{Y} \times \mathbf{\Sigma} \times \mathbf{Q} \to \mathbf{Y}$ with the following properties:*

1) *The state evolution function, $\mathcal{G}$, is given by equations* (15) *and* (19)*. This function determines the evolution of the vector of beliefs, P, and the vector of wealth positions, W.*

2) *When the Arrow security prices are given by $Q(\sigma'|\sigma)$ and when $y' = \mathcal{G}(y)$ the implied consumption plan solves Problem 1.*

3) *The goods market clears for all σ' where $c_i(\sigma')$ solves Problem 1:*

$$\sum_{i=1}^{N} c_i(\sigma') = N\varepsilon. \tag{20}$$

4) *The Arrow securities markets clear for all σ' where $a_i(\sigma') = W_i(\sigma') - H(\sigma')$:*

$$\sum_{i=1}^{N} a_i(\sigma') = 0. \tag{21}$$

In Proposition 2, we show that, in equilibrium, human wealth is a number that does not depend on the state and we derive an expression for the equilibrium price function $Q(\sigma'|\sigma)$.

PROPOSITION 2: *In a recursive equilibrium:*

1) Individual human wealth is given by Eq. (22)*,*

$$H = \frac{\varepsilon}{1 - \beta(1 - \delta)}. \tag{22}$$

2) The price of an Arrow security is given by Eq. (23)*,*

$$Q(\sigma'|\sigma) = \beta p(\boldsymbol{x}')\frac{\sum_{i=1}^{N} \mathbb{P}_i(s')x_i' W_i(\sigma)}{N(\sigma')H}, \tag{23}$$

where $N(\sigma') = \sum_i x_i'$ *is the number of surviving agents at time* $t+1$ *and* $N(\sigma')H$ *is aggregate human wealth.*

For a proof of Proposition 2 see Appendix D. In the next section, we will show how the pricing function, $Q(\sigma'|\sigma)$, depends on the assumptions about the information structure and the number of agents.

IV. Equilibrium Behaviour Under Two Different Assumptions

Next, we study the evolution of asset prices and the wealth distribution under two different assumptions. First, in Section IV.A, we assume that $\mathbb{P}_i(s') = \mathbb{P}(s')$ for all i We call this the *common knowledge economy* and we refer to the outcome of this version of our model as a rational expectations equilibrium. In Section IV.B we allow beliefs to differ and we ask and answer the question: Do markets reveal enough information for the economy to converge to a rational expectations equilibrium? We call this the *heterogeneous beliefs economy.*

A. The Common Knowledge Economy

When beliefs about the probability of s' are common, we can factor out $\mathbb{P}_i(s')$ from the sum in Eq. (23) and write the expressions for $Q(\sigma'|\sigma)$ as follows,

$$Q(\sigma'|\sigma) = \beta\mathbb{P}(s')p(\boldsymbol{x}')\theta(\boldsymbol{x}'), \tag{24}$$

where, using $\sum_{i=1}^{N} x_i' = N(\sigma')$,

$$\theta(\boldsymbol{x}') = 1 + \frac{\sum_{i=1}^{N} a_i(\sigma)x_i'}{N(\sigma')H}.$$

The term $\theta(\boldsymbol{x}')$ corrects Arrow security prices for mortality risk and we need to keep track of this term in our simulations to ensure that asset markets clear. This term disappears in the large N limit because each cohort is perfectly insured. As

$N \to \infty$, $\theta(\boldsymbol{x}') \to 1$ and we obtain the limiting expression[10]

$$Q(\sigma'|\sigma) = \beta \mathbb{P}(s') p(\boldsymbol{x}'). \tag{25}$$

Consider next how the wealth distribution evolves over time. The wealth evolution equation is given by Eq. (19), which we reproduce below,

$$W_i(\sigma') = x_i' \mathbb{P}(s') p(\boldsymbol{x}') \left[\frac{\beta W_i(\sigma)}{Q(\sigma'|\sigma)} \right] + (1 - x_i')H. \tag{26}$$

Combining this with Eq. (24) and using the fact that H is state-independent gives

$$W_i(\sigma') = x_i' W_i(\sigma) + (1 - x_i')H. \tag{27}$$

In the large N limit, there is no aggregate mortality risk and, in this case, we obtain an equivalent expression for $W_i(s)$

$$W_i(s') = x_i' W_i(s) + (1 - x_i')H. \tag{28}$$

Eq. (28) implies that in the large N economy, the wealth of the person with index i, contingent on her survival, is time invariant.

In a finite population, the variable $\theta(\boldsymbol{x}')$ plays a non-trivial role. Suppose, for example, that in period 1 there are two people. One person has positive financial assets equal to a and the other has negative financial assets equal to $-a$. In that economy, the rich person consumes more than the poor person for as long as they are both alive. But if one person dies and is replaced by a new person with wealth H, all debts are cancelled and the economy enters an absorbing state with an egalitarian wealth distribution. The wealth reallocation that occurs as a consequence of mortality risk is encoded into the random variable $\theta(\boldsymbol{x}')$.

B. *The Heterogeneous Belief Economy*

Next, we turn to the case where people have different beliefs. In this case, $\mathbb{P}_i(s')$ can no longer be factored out of the summation in Eq. (23) and instead of Eq. (24) we obtain the following expression for the price of an Arrow security,

$$Q(\sigma'|\sigma) = \beta p(\boldsymbol{x}') \left(\frac{\sum_{i=1}^N \mathbb{P}_i(s') W_i(\sigma) x_i'}{N(\sigma')H} \right) \equiv \beta \mathbb{P}_{\text{imp}}(\sigma') p(\boldsymbol{x}'), \tag{29}$$

where $\mathbb{P}_{\text{imp}}(\sigma')$ is the probability of state σ' that would be inferred from market prices if market participants believed that they were living in a common knowledge economy. We henceforth refer to $\mathbb{P}_{\text{imp}}(\sigma')$ as the *implied probability*.

[10]Notice that $\text{plim}_{N\to\infty} N^{-1} \sum_i a_i(\sigma) x_i' = 0$, using market clearing and assuming that $\text{plim}_{N\to\infty} N^{-2} \sum_i a_i^2 = 0$, which turns out to be true provided δ remains fixed as $N \to \infty$. Hence $\text{plim}_{N\to\infty} \theta(\boldsymbol{x}') = 1$.

Now, since who dies and who survives is independent from both wealth and beliefs one has, in the large N limit,[11]

$$\operatorname*{plim}_{N\to\infty} \frac{\sum_{i=1}^{N} \mathbb{P}_i(s')W_i(\sigma)x_i'}{N(\sigma')H} = \operatorname*{plim}_{N\to\infty} \left(\frac{\sum_{i=1}^{N} \mathbb{P}_i(s')W_i(s)}{NH} \right) \tag{30}$$

In the large N limit, the implied probability only depends on the future realisation of s but not on the mortality state. It is given by

$$\mathbb{P}_{\text{imp}}(s') \equiv \frac{\sum_{i=1}^{N} \mathbb{P}_i(s')W_i(s)}{NH}. \tag{31}$$

It is the wealth weighted average probability, and therefore differs from the true probability, $\mathbb{P}(s')$, which is the flat average of individual subjective probabilities.

Using the definition of $\mathbb{P}_{\text{imp}}(s')$, the analogue of Eq. (28) for the heterogeneous belief case is given by Eq. (32),

$$W_i'(s') = x_i' \frac{\mathbb{P}_i(s')}{\mathbb{P}_{\text{imp}}(s')} W_i(s) + (1 - x_i')H. \tag{32}$$

As we explain in Section V.A, $\mathbb{P}_i$ and W_i are strongly coupled by the dynamics, Eq. (32). Therefore, we cannot split $\mathbb{P}_{\text{imp}}(s')$ into the product of $\mathbb{P}(s')$ and $\operatorname{plim}_{N\to\infty} \sum_i \{W_i(\sigma)/N\}$ as we did in the common knowledge economy. This failure of independence generates fat-tails in the wealth distribution and it implies that the implied probability, $\mathbb{P}_{\text{imp}}(s')$, and the true probability, $\mathbb{P}(s')$, can differ *even in the large N limit*.

In Section V.D we will use these expressions to study the implications of our self-referential economy for the prices of debt and equity.

V. Results from Simulated Data

In Section V we illustrate the properties of our model by reporting some statistics for simulated data in a calibrated version. In subsection V.A we report the results of these simulations and in subsection V.C we derive the properties of some statistics for the large N continuous time limit.

A. A numerical simulation

We simulated an economy with one million agents for 300 years and we chose the period length to be one week. We normalized the weekly endowment to 1 and

[11] We use here the fact that if η_i and ξ_i are independent random variables, then

$$\operatorname*{plim}_{N\to\infty} N^{-1} \sum_{i=1}^{N} \eta_i \xi_i = \operatorname*{plim}_{N\to\infty} \left(N^{-1} \sum_{i=1}^{N} \eta_i \right) \left(N^{-1} \sum_{i=1}^{N} \xi_i \right),$$

and choose $\eta_i = \mathbb{P}_i W_i$ and $\xi_i = x_i'$.

we chose the annual discount rate to be 0.97 which corresponds to an equilibrium annual real interest rate, in an endowment economy, of 3%. These are relatively uncontroversial choices. In Figure 1 we graph some data from a single simulation of this calibrated version of our model.

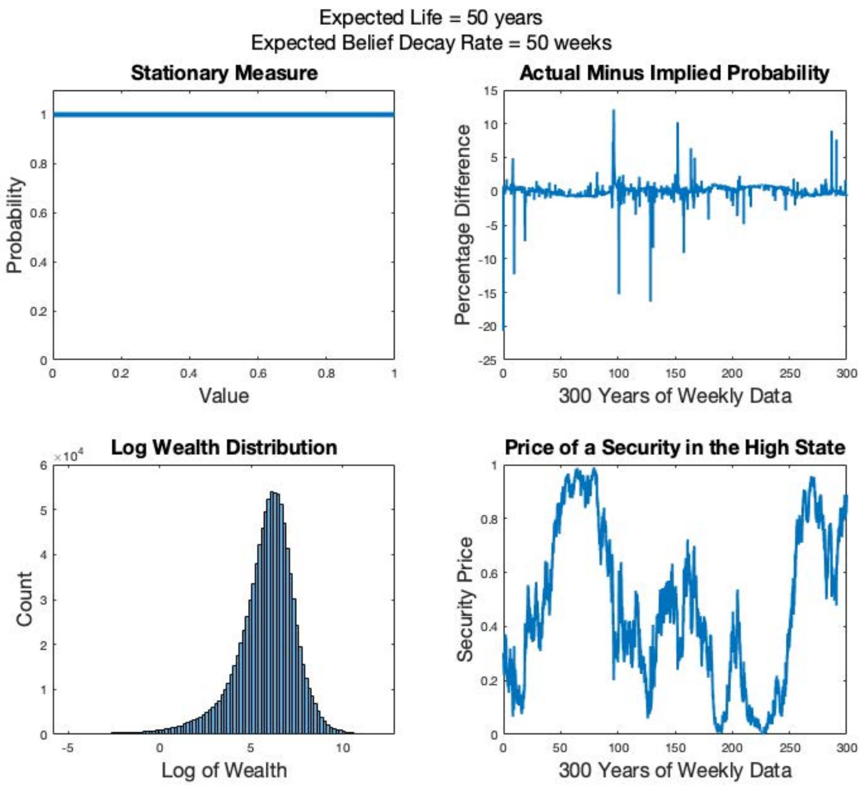

Figure 1. 300 Years of Simulated Weekly Data in an Economy with One Million People

The properties of our simulations are sensitive to two key parameters. The first is δ which governs the replacement rate of new agents. The second is λ, which governs the weight that agents place on new information in updating their beliefs. We chose $\delta = 3.9 \times 10^{-4}$ which gives a life expectancy for the average agent of 50 years for the chosen time step of one week. In our model, life expectancy is independent of age and our choice for δ is consistent with US life expectancy tables for which a crude age-weighted average of survival probabilities delivers a number close to 50 years.

We chose a value for λ equal to $\delta^{1/2}$, i.e $\alpha = 1$. For our calibrated value of δ, this choice implies that it takes 50 weeks (approximately one year) for the effect of the prior to be swamped by new data. We constrained δ and λ to be linked in this way because, in the large N continuous time limit, this choice of parameters implies that $\mathcal{P}_\infty(\mathbb{P})$ is uniform on $[0,1]$ – cf. Eq. (5).[12]

The top left panel of Figure 1 graphs the invariant measure $\mathcal{P}_\infty(\mathbb{P})$. The other three panels present some key data for a single simulation of 300 years of weekly data. The top right panel is the percentage difference between $\mathbb{P}(s')$ and $\mathbb{P}_{\text{imp}}(s')$. This difference is a measure of how wrong the market can be as a measure of the true probability. For much of the sample this difference is less that 1% but there are times when this deviation exceeds $\pm 15\%$. Such large discrepancies are quite remarkable in view of the size of the market (one million participants) and are the consequence of the emergent wealth inequalities in our model.

The bottom right panel shows the time series behaviour of the price for delivery of a commodity in the high state. This price wanders randomly over the interval $[0,1]$ and sometimes it moves substantially in a short period of time. The bottom left panel is the log of the wealth distribution. In the following subsection we explore the properties of this distribution and we show that it shares many characteristics in common with empirical wealth distributions in Western economies.

B. *Exploring the Empirical Wealth Distribution*

When we embed our learning mechanism in a market economy, a somewhat unexpected effect appears. While our model is constructed in such a way that no agent is better informed than any other, some agents are by chance, and temporarily, much more successful than others. This allows these agents to accumulate wealth through the multiplicative process described in Eq. (32), reproduced below

$$W_i'(s') = x_i' \frac{\mathbb{P}_i(s')}{\mathbb{P}_{\text{imp}}(s')} W_i(s) + (1 - x_i')H. \tag{32}$$

Multiplicative wealth processes of this form are well-known to generate important wealth inequalities, as we explain in Section V.C. In Figure 2 we graph the Lorenz curve for the time average of 250 equally spaced samples of the wealth distribution in our simulated data. The Lorenz curve is a graphical representation of inequality which plots the cumulative percentage of wealth on the y-axis against the percentile of the population on the x-axis. One popular index of inequality is the Gini coefficient which is equal to twice the area between the 45 degree line and the Lorenz curve.

For our numerical data, the Gini coefficient is equal to 0.7. A value of 0 would represent a completely equal distribution and a value of 1 would represent a

[12]Our parameterisation is consistent with the fact that although many financial time series appear to be non -stationary, they do not appear to pile up around a small number of values. A value of λ^2 that was greater than δ would imply that $\mathbb{P}$ piles up around 1/2. A value less than δ would cause $\mathbb{P}$ to spend a lot of time around $\mathbb{P} = 0$ or to jump occasionally and spend a similar amount of time around $\mathbb{P} = 1$. Because we do not observe either of those features in data we chose a value of $\lambda^2 = \delta$ which implies that $\mathbb{P}$ wanders randomly in $[0,1]$.

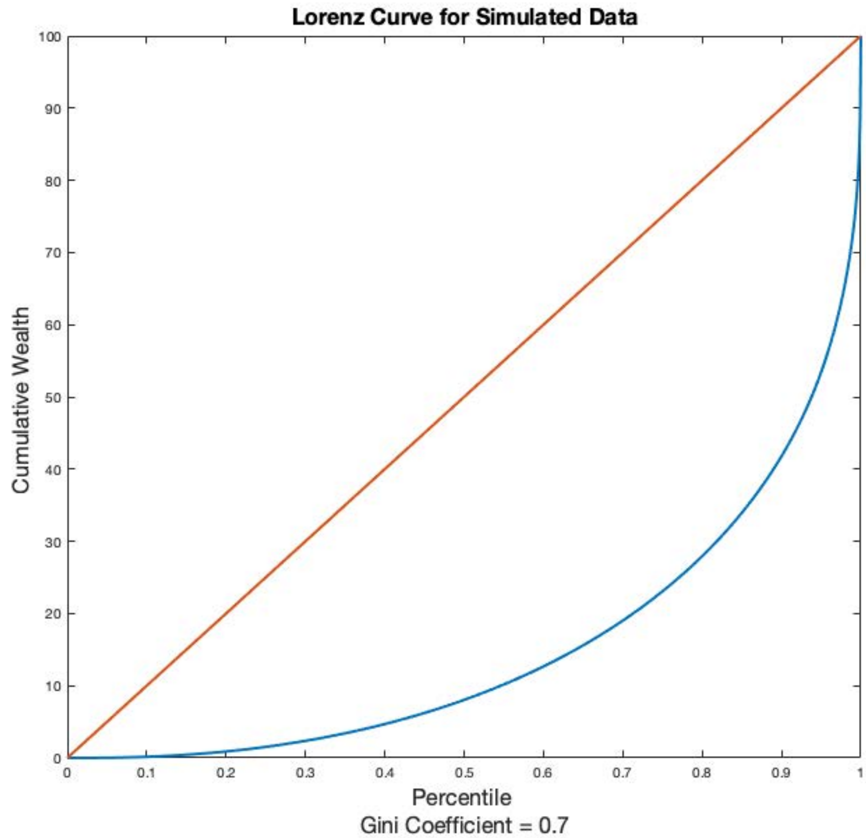

Figure 2. The Lorenz Curve for a Single Simulation

distribution where one person owns everything. As we show below a value of 0.7 is close to observed Gini coefficients for the wealth distribution in the data.

Table 1 reports data from a selection of countries. This table shows that a Gini coefficient of 0.7 is well within the bounds of empirical data which varies between a low of 0.55 for China in 2008 and a high of 0.85 for the United States in 2019. To explore the nature of the wealth distribution further we define $F(W)$ to be the cumulative distribution function (cdf) of wealth and define $G(W) \equiv 1 - F(W)$ to be the complementary cdf. In Figure 3, we plot $\log G(W)$ against $\log(W)$ for values of $\log(W)$ greater than zero. This figure reveals a power-law tail of the form $G(W) \sim W^{-\mu}$, and a regression of $\log(G(W))$ on $\log(W)$ for the linear portion of the plot provides an estimate of the tail index of $\mu = 1.4$. Note that $G(W) \sim W^{-\mu}$

[13]Wikipedia https://en.wikipedia.org/wiki/List_of_countries_by_wealth_equality Retrieved December 6'th 2020.

Country	2008	2019
China	0.55	0.7
United Kingdom	0.7	0.75
Italy	0.7	0.77
France	0.73	0.7
Switzerland	0.74	0.87
United States	0.8	0.85

Table 1—Wealth Ginis' For a Selection of Countries in 2008 and 2019[13]

corresponds to a probability distribution function (pdf) $\varrho(W) \sim W^{-1-\mu}$. A person who is neither a borrower nor a lender has zero financial assets and her net worth would be equal to the discounted present value of her labour income. For our calibration, this number, which we refer to as human wealth, is equal to 1,032 weeks of income.[14] In the common knowledge economy, the wealth distribution would be egalitarian, the Gini coefficient would be 0 and everyone would have wealth equal to H. Instead, in our economy, there is considerable inequality.

A person at the 50'th percentile of the wealth distribution is a net borrower who has total wealth equal to 39% of human wealth. In contrast, a person at the $99'th$ percentile in the wealth distribution has total wealth equal to 892% of human wealth and the person at $99.9'th$ percentile has total wealth of $4,999\%$. Wealth becomes highly concentrated because market prices do not reflect average beliefs. Instead they reflect wealth weighted beliefs. In equilibrium, wealth and market prices are correlated in a way that leads to a self-reinforcing mechanism whereby a few people, by chance, get lucky and become very rich.

C. The Behaviour of Wealth in the Large N Limit

We can learn quite a bit about the dynamics of wealth by analyzing the properties of Eq. (32). Using this equation, one may derive the following expression for the average return for agent i between dates t and $t+1$, conditional on surviving:[15]

$$R_i \equiv \mathbb{E}_t \left[\frac{W_i'}{W_i} - 1 \right] = \frac{(\mathbb{P} - \mathbb{P}_{\text{imp}})(\mathbb{P}_i - \mathbb{P}_{\text{imp}})}{\mathbb{P}_{\text{imp}}(1 - \mathbb{P}_{\text{imp}})}. \tag{33}$$

[14]Human wealth is defined by the expression $H = 1/(1 - \beta(1-\delta))$. For our calibration the weekly discount rate is $0.97^{1/52}$ and the survival probability, $(1-\delta)$, is equal to 3.9×10^{-4}. This leads to a value of $H = 1,031$ measured in weeks of income.

[15]Eq. (33) follows since

$$\begin{aligned} R_i \equiv \mathbb{E}_t \left[\frac{W_i'}{W_i} - 1 \right] &= \left[\mathbb{P} \times \frac{\mathbb{P}_i}{\mathbb{P}_{\text{imp}}} + (1-\mathbb{P}) \times \frac{1-\mathbb{P}_i}{1-\mathbb{P}_{\text{imp}}} - 1 \right] \\ &= \frac{(\mathbb{P} - \mathbb{P}_{\text{imp}})(\mathbb{P}_i - \mathbb{P}_{\text{imp}})}{\mathbb{P}_{\text{imp}}(1 - \mathbb{P}_{\text{imp}})}. \end{aligned}$$

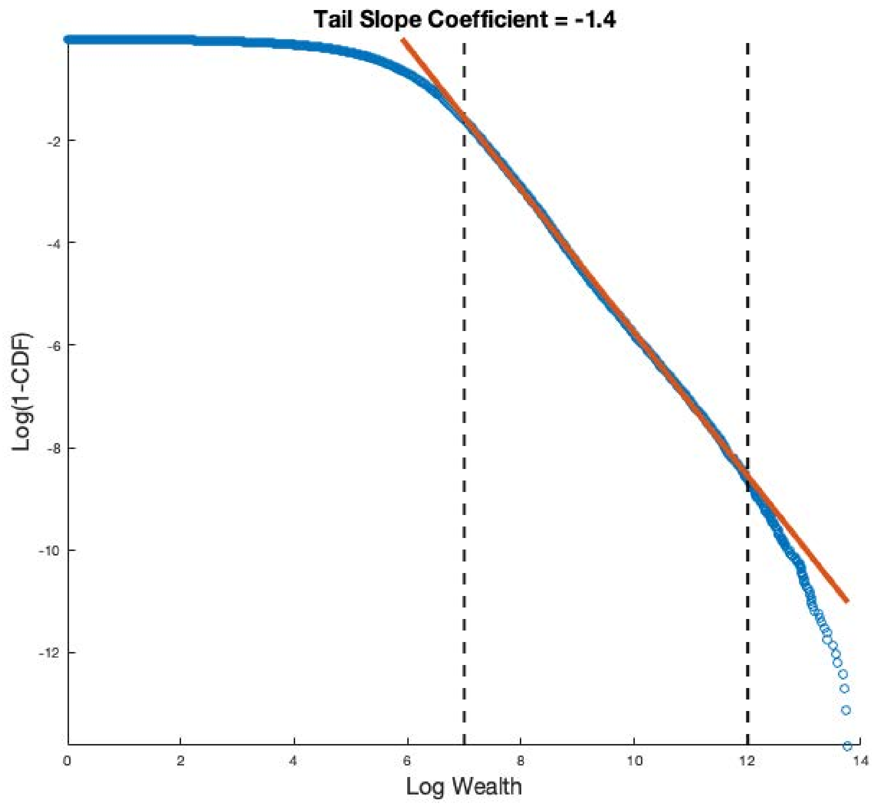

Figure 3. Estimate of the Tail Parameter in 1,000 Years of Simulated Monthly Data

Several interesting conclusions can be drawn from Eq. (33). First, in the common knowledge economy where $\mathbb{P} \equiv \mathbb{P}_{\text{imp}}$, agents cannot expect to make money on average, even temporarily.

Second, when an agent's belief $\mathbb{P}_i$ is larger than the market probability $\mathbb{P}_{\text{imp}}$, her expected gain is positive if the actual probability $\mathbb{P}$ is also greater than $\mathbb{P}_{\text{imp}}$, and negative otherwise.

Finally, since agents are assumed to act on the assumption that their estimate of the probability is an unbiased estimate of the true probability, they also believe that their trades will be profitable on average and proportional to $(\mathbb{P}_i - \mathbb{P}_{\text{imp}})^2$. In other words, they expect to make a larger profit, the further is their belief from the probability implied by the market price. This implies that there is no incentive for agents to align their beliefs with the observable implied probability, since this would reduce their subjective expected profit. Everybody in this economy, believes that they know more than the market.

In Eq. (34) we derive an expression for the average of the square of the relative change of wealth for surviving agents:[16]

$$\mathbb{E}_t\left[\left(\frac{W_i'}{W_i}-1\right)^2\right]=\frac{\left(\mathbb{P}-2\mathbb{P}\mathbb{P}_{\text{imp}}+\mathbb{P}_{\text{imp}}^2\right)(\mathbb{P}_i-\mathbb{P}_{\text{imp}})^2}{\mathbb{P}_{\text{imp}}^2(1-\mathbb{P}_{\text{imp}})^2}. \tag{34}$$

One sees from this equation that "bold beliefs", corresponding to a large difference between $\mathbb{P}_i$ and the market probability $\mathbb{P}_{\text{imp}}$, leads to a larger variance of gains. Eq. (34) explains why our model generates large wealth inequalities. For surviving agents, the wealth dynamic is a multiplicative random process with a time dependent and agent dependent variance. This multiplicative process is reset to 1 at a Poisson rate δ, i.e. when an agent dies.

Multiplicative random process with reset have been widely studied in the literature (see e.g. Kesten (1973); Bouchaud and Mézard (2000); Benhabib et al. (2018); Gabaix (2009); Benhabib and Bisin (2011); Gabaix et al. (2016)) and it is known that such processes lead to a stationary distribution with a power-law tail with a pdf $\varrho(W)$ and a complementary cdf $G(W)$ of the form,

$$\varrho(W)\sim_{W\to\infty}W^{-1-\mu},\qquad G(W)\sim_{W\to\infty}W^{-\mu}, \tag{35}$$

where the exponent μ depends on the parameters of the problem. We discuss in Appendix E how μ can be approximately computed. We find, in particular, that $\mu>1$ whenever $\delta>0$.

Random variables that behave 'like' the Pareto distribution for large W are said to possess a 'Pareto tail'. These distributions can be sorted into three classes depending on the value of the tail parameter μ. A Pareto-tailed distribution is well defined for all positive μ but when $0<\mu\leq 1$, the mean and all higher moments do not exist. When $1<\mu\leq 2$, the mean exists but the variance and higher moments do not exist and for $\mu>2$, the distribution has a finite mean and a finite variance. In our example, as in the data, we find a value of μ between 1 and 2 which implies that the wealth distribution has a finite first moment but all higher order moments are not well defined.

In conclusion, wealth inequalities in our model arise from the multiplicative nature of wealth dynamics which makes successful bold bets highly profitable. The flip side of this statement is that unsuccessful bold bets are ruinous and lead the person who makes such bets into poverty. People who agree with the market

[16] Eq. (34) follows since

$$\begin{aligned}\mathbb{E}\left[\left(\frac{W_i'}{W_i}-1\right)^2\right]&=\left[\mathbb{P}\left(\frac{\mathbb{P}_i}{\mathbb{P}_{\text{imp}}}-1\right)^2+(1-\mathbb{P})\left(\frac{1-\mathbb{P}_i}{1-\mathbb{P}_{\text{imp}}}-1\right)^2\right]\\&=\frac{\left(\mathbb{P}-2\mathbb{P}\mathbb{P}_{\text{imp}}+\mathbb{P}_{\text{imp}}^2\right)(\mathbb{P}_i-\mathbb{P}_{\text{imp}})^2}{\mathbb{P}_{\text{imp}}^2(1-\mathbb{P}_{\text{imp}})^2}.\end{aligned}$$

.

belief have a low expected subjective gain from trading. People who disagree may either become spectacularly rich, or spectacularly poor.

D. Debt and Equity in the Heterogeneous Belief Economy

In this section, we flesh out the idea that s' is a signal, transmitted through a social network, of the public's commonly held belief that the stock market will take on a high value, rather than a low value. We assume that N is large and we derive formulae that hold exactly in the large N limit.

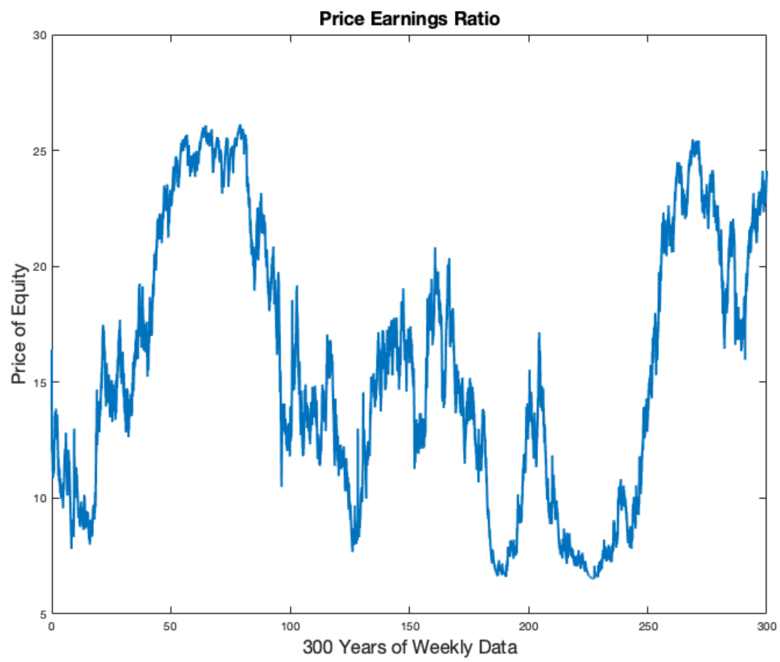

FIGURE 4. THE BEHAVIOUR OF THE PRICE OF EQUITY IN 300 YEARS OF SIMULATED MONTHLY DATA

We derive trading rules for traders who buy and sell Arrow securities. But there is no reason to restrict ourselves to securities of this kind and the same equilibrium we have described above can be supported by any set of securities with independent payoffs that span the space of possible outcomes. Here, we show that the equilibrium can be supported by a security that pays one commodity in both states; we call this security *debt*, and a security that pays d units if $s = \{1\}$ and zero otherwise. We call this security *equity*.

PROPOSITION 3: *For the large N economy, equilibrium can be supported by trades in debt and equity. Debt is a security that costs Q units of commodities* *in state s and pays 1 commodity in state s'. Equity is a security that costs $p_E(s)$ units of commodities in state s and pays $p_E(s') + d$ in state $s' = \{1\}$ and $p_E(s')$ in state $s' = \{0\}$, where*

$$p_E(s) = \frac{d\beta}{2}\left[\frac{2\mathbb{P}_{imp} - 1}{1 - \beta(1-\delta)} + \frac{1}{1-\beta}\right], \tag{36}$$

$$Q = \beta. \tag{37}$$

For a proof of Proposition 3 see Appendix F. In Figure 4 we have graphed the value of $p_{\mathrm{E}}(s)$ for the data simulated in Figure 1. To compute this series we normalized the dividend payment to 1/52 to make the units comparable to an expected weekly dividend payment. This series has many characteristics in common with the price dividend ratio in US data for realized values of the S&P.

In the introduction to this paper, we promised to develop a model in three stages. In stage 1, we presented a model where people disagree about the value of a public signal, but where that signal has no economic consequences. In stage 2, we allowed people to trade Arrow securities, conditional on that signal, but we did not explain how Arrow securities are related to real world trades in debt and equity. In stage 3, discussed in this section, we have shown that, in the large N limit, trades in a pair of Arrow securities can be replicated by trades in two securities that resemble debt and equity.

In the world where agents trade debt and equity, firms pay dividends contingent on the realization of a public signal. As we suggested in the introduction, one possible interpretation of this signal is the publication of an optimistic article by a journalist for a national newspaper which triggers a dividend payout by firms. The fact that firms react to the signal by making a dividend payment in good states, but not in bad states, is a self-fulfilling action which confirms the beliefs of market traders that they should be willing to pay more for shares in the firm.

VI. Conclusion

We have a constructed a theory of beliefs in which people exchange information through both market and non-market interactions. Non-market interaction in social networks generates an aggregate signal which reflects average public opinion. Market exchange through the purchase and sale of financial assets allows people to bet on their beliefs. Importantly, market prices reveal information about wealth weighted beliefs but it is flat weighted beliefs which generate the public signal.

One is led to the question: Why do people continue to bet with each other when these bets are highly risky? The answer we propose is that everyone in our economy thinks that the market is wrong and that by betting, they will be able to make money on average. They do not use the implied probability revealed by the markets to improve their estimate of $\mathbb{P}$, since this trading strategy is, in their opinion, sub-optimal. Quite remarkably, the coupled dynamics of individual wealth and beliefs leads to a fat-tailed distribution of wealth that prevents markets to faithfully reflect agents' beliefs, therefore preventing agents to learn the true value of $\mathbb{P}$.[17]

Why are there no Warren Buffets who invest for the long run by guessing that the probability of a successful outcome will be equal to the mean of the invariant distribution? Our answer is that for any reasonable time period, the future value of s is far better approximated from averaging the frequency of its recent realizations than by assuming that it is drawn form the unconditional long-run measure. In the long run the expected probability that $s' = \{1\}$ is 1/2, but, as Keynes famously quipped: "In the long run we are all dead". We believe that our quasi non-ergodic model aptly illustrates what Keynes had in mind.

REFERENCES

Klaus Adam, Albert Marcet, and Juan Pablo Nicolini. Stock market volatility and learning. *Journal of Finance*, 71(1):33–82, 2016.

Klaus Adam, Albert Marcet, and Johannes Beute. Stock price booms and expected capital gains. *American Economic Review*, 107(8):2352–2408, 2017.

Philip W. Anderson. Spin glass vi: Spin glass as cornucopia. *Physics Today*, 42 (9):9–11, 1989.

Kenneth J. Arrow. The role of securities in the optimal allocation of risk bearing. *Review of Economic Studies*, 31(2):91–96, 1964.

Alejandro Badel, Moira Daly, Mark Huggett, and Martin Nybom. Top earners: Cross-country facts. *Institute for Evaluation of Labour Market and Education Policy Working Paper 2017:9*, 2017.

Pablo Beker and Emilio Espino. The dynamics of efficient asset trading with heterogeneous beliefs. *Journal of Economic Theory*, 166:189–229, 2011.

Jess Benhabib and Alberto Bisin. The distribution of wealth and fiscal policy in economies with finitely lived agents. *Econometrica*, 79(1):123–157, 2011.

Jess Benhabib, Alberto Bisin, and Shenghao Zhu. Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, 56(4):1261–91, 2018.

Olivier J. Blanchard. Debt, deficits, and finite horizons. *Journal of Political Economy*, 93(April):223–247, 1985.

Lawrence Blume and David Easley. If you're so smart, why aren't you rich? *Econometrica*, 74(4):929–966, 2006.

[17]Introducing a wealth tax, or an inheritance tax, tends to reduce inequalities and helps prices reveal private beliefs – see Appendix E.

Jean-Philippe Bouchaud and Marc Mézard. Wealth condensation in a simple model of the economy. *Physica A: Statistical Mechanics and its Applications*, 282(3-4):536–545, 2000.

William A. Brock and Steven N. Durlauf. Discrete Choice with Social Interactions. *The Review of Economic Studies*, 68(2):235–260, 04 2001.

Timothy Cogley and Thomas J. Sargent. The market price of risk and the equity premium: A legacy of the great depression? *Journal of Monetary Economics*, 5(3):454–476, April 2008.

Timothy Cogley and Thomas J. Sargent. Diverse beliefs, survival and the market price of risk. *Economic Journal*, 119:354–376, March 2009.

Pablo G. Debenedetti and Frank H. Stillinger. Supercooled liquids and the glass transition. *Nature*, 410(6825):259–267, 2001. URL `http://dx.doi.org/10.1038/35065704`.

Roger E. A. Farmer. *The Macroeconomics of Self-Fulfilling Prophecies.* MIT Press, Cambridge, MA, second edition, 1999.

Roger E. A. Farmer. Pricing assets in a perpetual youth model. *Review of Economic Dynamics*, 30:106–124, 2018.

Roger E. A. Farmer. The indeterminacy school in macroeconomics. *Oxford Research Encyclopedia of Economics and Finance*, online publication, April 2020.

Roger E. A. Farmer. The importance of beliefs in shaping macroeconomic outcomes. *Oxford Review of Economic Policy*, forthcoming, 2021.

Roger E. A. Farmer, Carine Nourry, and Alain Venditti. Debt deficits and finite horizons, the stochastic case. *Economics Letters*, 111:47–49, 2011.

Xavier Gabaix. Power laws in economics and finance. *Annual Review of Economics*, 1(1):255–294, 2009.

Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lyons, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84(6):2071–2111, 2016.

Peter Hánggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after Kramers. *Reviews of modern physics*, 62:251–341, 1990.

Ulrich Horst. *Ergodicity and Nonergodicity in Economics*, pages 1–6. Palgrave Macmillan UK, London, 2017.

Harry Kesten. Random difference equations and renewal theory for products of random matrices. *Acta Math.*, 131:207–248, 1973. doi: 10.1007/BF02392040. URL `https://doi.org/10.1007/BF02392040`.

John Maynard Keynes. *The General Theory of Employment, Interest and Money.* MacMillan and Co., London and Basingstoke, 1936. 1973 edition published for the Royal Economic Society, Cambridge.

Alan Kirman. Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108:137–156, 1993.

José Moran, Antoine Fosset, Michael Benzaquen, and Jean-Philippe Bouchaud. Schrodinger's ants: a continuous description of kirman's recruitment model. *Journal of Physics: Complexity*, 1(3):035002, Aug 2020a.

José Moran, Antoine Fosset, Davide Luzzati, Jean-Philippe Bouchaud, and Michael Benzaquen. By force of habit: Self-trapping in a dynamical utility landscape. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(5): 053123, 2020b. doi: 10.1063/5.0009518.

Patrick A.P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54:60–71, 1958.

Federico Guglielmo Morelli, Michael Benzaquen, Marco Tarzia, and Jean-Philippe Bouchaud. Confidence collapse in a multihousehold, self-reflexive dsge model. *Proceedings of the National Academy of Sciences*, 117(17): 9244–9249, 2020. ISSN 0027-8424. doi: 10.1073/pnas.1912280117. URL `https://www.pnas.org/content/117/17/9244`.

Robin Pemantle. A survey of random processes with reinforcement. *Probability surveys*, 4:1–79, 2007.

Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15(12):1261–1221, 2019.

Thomas Piketty and Gabriel Zucman. Capital is Back: Wealth-Income Ratios in Rich Countries 1700–2010 . *The Quarterly Journal of Economics*, 129(3): 1255–1310, 05 2014.

Nancy L. Stokey, Robert E. Lucas, Jr., and with Edward C. Prescott. *Recursive Methods in Economic Dynamics.* Harvard University Press, Cambridge, MA, 1989.

Appendix A: The continuous time limit

A1. Derivation of Eq. (5)

Introducing a change of variable u such that $\mathbb{P} = \frac{1}{2} + u$, one can convert Eq. (4) into:

$$(1-\hat{\delta})^2 \mathcal{P}_{t+1}(u) = \frac{1-\hat{\delta}-\hat{\lambda}}{2}\left[\mathcal{P}_t\left(\frac{u-\hat{\lambda}/2}{1-\hat{\delta}}\right) + \mathcal{P}_t\left(\frac{u+\hat{\lambda}/2}{1-\hat{\delta}}\right)\right] + u\left[\mathcal{P}_t\left(\frac{u-\hat{\lambda}/2}{1-\hat{\delta}}\right) - \mathcal{P}_t\left(\frac{u+\hat{\lambda}/2}{1-\hat{\delta}}\right)\right] \tag{A1}$$

where $\hat{\lambda} = \lambda(1-\delta)$ and $\hat{\delta} = \delta + \hat{\lambda}$.

In the following analysis we assume long memory ($\lambda \ll 1$) and rare mistakes ($\delta \ll 1$) by focusing on the limit where $\lambda, \delta \to 0$ with $\delta = \alpha\lambda^2$ for fixed $\alpha = O(1)$. Expanding Eq. (A1) to order λ^3 yields:

$$\Delta_t = \delta\left[u\mathcal{Q}\right]' + \frac{\lambda^2}{2}\left[(\frac{1}{4} - u^2)\mathcal{Q}\right]'' - 2\lambda\delta\left[u^2\mathcal{Q}\right]'' - \frac{\lambda^3}{2}\left[(\frac{u}{12} - \frac{u^3}{3})\mathcal{Q}'' - u^2\mathcal{Q}' + \frac{5}{12}\mathcal{Q}'\right]' + O(\lambda^4), \tag{A2}$$

where primes denote derivatives with respect to u, $\mathcal{P}(u) \equiv (1-\hat{\delta})\mathcal{Q}(u(1-\hat{\delta}))$, and $\Delta_t \equiv \mathcal{Q}_{t+1}(u) - \mathcal{Q}_t(u)$. Note that the last two terms of Eq. (A2) are of order λ^3, and we will neglect them in the following approximation.

In the small δ, λ limit, Eq. (A2) converges to the following continuous time Fokker-Planck equation for $\mathcal{P}$:

$$\frac{1}{\lambda^2}\frac{\partial\mathcal{P}}{\partial t} = \alpha\left[u\mathcal{P}\right]' + \frac{1}{2}\left[(\frac{1}{4} - u^2)\mathcal{P}\right]''. \tag{A3}$$

This equation coincides with the continuous time description of Kirman's ant recruitment model (Kirman, 1993), for which a lot is known (see Moran (1958) for recent results and references).

In particular the stationary distribution $\mathcal{P}^*$ is is described by the following second order differential equation.

$$\alpha\left[u\mathcal{P}^*\right]' + \frac{1}{2}\left[(\frac{1}{4} - u^2)\mathcal{P}^*\right]'' = 0. \tag{A4}$$

The solution to this equation is given by

$$\mathcal{P}_\infty(u) = \frac{\Gamma(2\alpha)}{\Gamma^2(\alpha)} \left(\frac{1}{4} - u^2\right)^{\alpha-1}, \tag{A5}$$

which corresponds to Eq. (5) in the text.

A2. Generalisation: non-linear feedback

The Fokker-Planck equation Eq. (A3) corresponds to the following stochastic differential equation:

$$\mathrm{d}\mathbb{P} = -\delta(\mathbb{P} - \frac{1}{2})\mathrm{d}t + \lambda\sqrt{\mathbb{P}(1-\mathbb{P})}\mathrm{d}W_t, \tag{A6}$$

where W_t is a Wiener noise. More generally, one can consider a sigmoidal feedback term $\mathcal{F}(\mathbb{P})$ mapping the average belief onto the true probability,

$$\mathbb{P}_{t+1} = \mathcal{F}(\mathbb{P}_t) \tag{A7}$$

with $\mathcal{F}(\mathbb{P}) = \mathbb{P}$ throughout the main part of the paper and in section above. In this case, one obtains as a stochastic differential equation

$$\mathrm{d}\mathbb{P} = -\partial_{\mathbb{P}}\mathcal{V}(\mathbb{P})\mathrm{d}t + \lambda\sqrt{\mathbb{P}(1-\mathbb{P})}\mathrm{d}W_t, \tag{A8}$$

where we have introduced a "potential function" $\mathcal{V}(x)$ such that

$$\partial_x\mathcal{V}(x) := \delta(x - \frac{1}{2}) + \lambda(x - \mathcal{F}(x)). \tag{A9}$$

For definiteness, consider a sigmoidal function $\mathcal{F}(x)$ defined as:

$$\mathcal{F}(x) = \frac{1}{2}\left(1 + \tan[\beta(x - \frac{1}{2})]\right) \tag{A10}$$

The corresponding potential $\mathcal{V}(x)$ is then given by

$$\mathcal{V}(x) = \frac{1}{2}(\delta+\lambda)u^2 - \frac{\lambda}{2\beta}\log\cosh\beta u; \qquad u := x - \frac{1}{2} \tag{A11}$$

For small β, $\mathcal{V}(x)$ has a unique minimum corresponding to $x = 1/2$. For $\beta > \beta_c = 2(1+\delta/\lambda)$, $\mathcal{V}(x)$ has two minima $x^* < 1/2$ and $1 - x^* > 1/2$ and one maximum at $x = 1/2$.

In the absence of the Wiener noise term, the dynamics of x would just be "rolling down" the potential slopes, selecting one of the minima of $\mathcal{V}(x)$ (corresponding to the stable solutions of $\mathcal{F}(x) = x$).

In the presence of noise and for $\beta > \beta_c$, the dynamics becomes a succession of long phases where $\mathbb{P}_t$ remains close to either x^* or $1-x^*$, separated by rapid switches from one minimum to the other. The time $\tau_\times$ needed to "climb up the hill" separating the two minima can be however very long when $\lambda \to 0$.

In fact, this time can be rather accurately computed by changing variables from $\mathbb{P}$ to ϕ where $\mathbb{P} = (1+\sin\phi)/2$, which allows one to get rid of the factor $\sqrt{\mathbb{P}(1-\mathbb{P})}$ in front of the Wiener noise, see e.g. Moran et al. (2020a). Using a standard approach (e.g. Hánggi et al. (1990)), one can then show that

$$\tau_\times \sim \lambda^{-1} e^{\Gamma/\lambda}, \qquad (\lambda \to 0),$$

where Γ can be fully computed (at least numerically) for *any* potential $\mathcal{V}(x)$. The exponential dependence of $\tau_\times$ in λ implies that (a) there is a strong separation of timescales in such models and (b) the precise value of $\tau_\times$ is unknowable in practice, as it is highly sensitive on the detailed value of the parameters of the model. Hence agents cannot be assumed to use the same learning rule. Since these switches can be interpreted as "crashes", the probability of such crashes is, in our simple model, unknowable much as the trajectories of a chaotic system are unknowable (for a related discussion, see Morelli et al. (2020)).

Appendix B: Dispersion of opinions

Taking the expectation of Eq. (6) over the realisation of s_t one gets:

$$\mathbb{E}[\mathbb{D}_{i,t+1}] = (1-\delta)\left[(1-\lambda)\mathbb{E}[\mathbb{D}_{i,t}] + \delta(\mathbb{P}_t - \frac{1}{2})\right] + \delta(1-\delta)\left[\frac{1}{2} - \mathbb{P}_t\right], \tag{B1}$$

or

$$\mathbb{E}[\mathbb{D}_{i,t+1}] = (1-\delta)(1-\lambda)\mathbb{E}[\mathbb{D}_{i,t}] \tag{B2}$$

which shows that $\mathbb{E}[\mathbb{D}_{i,t}]$ tends to zero when $t \to \infty$.

Now let us square Eq. (6) before taking the average over s_t. One now gets:

$$\begin{aligned} \mathbb{E}[\mathbb{D}^2_{i,t+1}] = (1-\delta)\left[(1-\lambda)^2\mathbb{E}[\mathbb{D}^2_{i,t}] + \delta^2\mathbb{E}[(\mathbb{P}_t - \frac{1}{2})^2]\right] \\ + \delta\left[\mathbb{E}[z^2] + \frac{\delta^2}{4} - \frac{\delta}{2} + (1-\delta)^2(1-\lambda^2)(\mathbb{P}_t^2 - \mathbb{P}_t)\right]. \end{aligned} \tag{B3}$$

Now taking further the expectation over the distribution $\mathcal{P}$ of the probability $\mathbb{P}$, and using

$$\mathbb{E}_{\mathcal{P}}[\mathbb{P}^2] = \frac{1+\alpha}{2(1+2\alpha)}, \qquad \alpha = \frac{\delta}{\lambda^2}, \tag{B4}$$

we obtain, in the limit $\delta, \lambda \to 0$, with α fixed,

$$\mathbb{E}^\star[\mathbb{D}_{i,t+1}^2] = (1-\delta)(1-\lambda)^2\mathbb{E}^\star[\mathbb{D}_{i,t}^2] + \frac{\delta}{6}\frac{2+\alpha}{1+2\alpha} + O(\delta^2), \tag{B5}$$

where $\mathbb{E}^\star$ means an expectation both over s and $\mathcal{P}$.

Hence in the stationary state where $\mathbb{E}^\star[\mathbb{D}_{i,t}^2]$ is independent of t one finds:

$$\mathbb{E}^\star[\mathbb{D}_i^2] \approx \frac{\delta}{6(1-(1-\delta)(1-\lambda)^2)}\frac{2+\alpha}{1+2\alpha}, \tag{B6}$$

and hence the result Eq. (7).

Appendix C: Solving the individual optimization problem

We conjecture that the value function has the form

$$A \log W_i(\sigma) + B, \tag{C1}$$

for unknown constants A and B. Substituting from Eq. (16) for $c_i(\sigma)$ in Eq. (14) and taking derivatives with respect to $W_i(\sigma')$ leads to the following Euler equation,

$$\frac{x_i(\sigma')Q(\sigma'|\sigma)}{c_i(\sigma)} = \frac{A\beta\mathbb{P}_i(\sigma')x_i(\sigma')}{W_i(\sigma')}, \tag{C2}$$

which holds state by state. Using the envelope condition $Ac_i(\sigma) = W_i(\sigma)$, which holds at every date and in every state, we can write Eq. (C2) as

$$x_i(\sigma')Q(\sigma'|\sigma)W_i(\sigma') = \beta\mathbb{P}_i(\sigma')x_i(\sigma')W_i(\sigma). \tag{C3}$$

Combining the budget equation, Eq. (16), which holds with equality with Eq. (C3) leads to the expression,

$$\sum_{\sigma'} \beta\mathbb{P}_i(\sigma')x_i(\sigma')W_i(\sigma) + \frac{W_i(\sigma)}{A} = W_i(\sigma). \tag{C4}$$

Because s' is independent of $\boldsymbol{x}'$

$$\sum_{\sigma'} \mathbb{P}_i(\sigma')x_i(\boldsymbol{x}') = \sum_{\boldsymbol{x}'} p(\boldsymbol{x}')x_i(\boldsymbol{x}') \sum_{s'} \mathbb{P}_i(s') = 1-\delta \tag{C5}$$

and thus by canceling terms and rearranging Eq. (C3) we arrive at the following value for A.

$$A = \frac{1}{1-\beta(1-\delta)} \tag{C6}$$

The constant B does not affect the solution and can be solved for by plugging the value of A into the expression

$$A\log(W_i) + B = \log\left(\frac{W_i}{A}\right) + \beta(1-\delta)\left[A\log(W_i) + B\right] \tag{C7}$$

and equating the coefficients on the constant terms.

It follows from Eq. (C3) that for all $x_i(\boldsymbol{x}') = 1$, that is, those who survive,

$$W_i(\sigma') = \beta\frac{\mathbb{P}_i(\sigma')}{Q(\sigma'|\sigma)}W_i(\sigma). \tag{C8}$$

This establishes the first term on the right side of Eq. (19). If $x_i(\boldsymbol{x}') = 0$ the newborn with index i has wealth H by assumption. This establishes the second term on the right side of Eq. (19).

Appendix D: Establishing the Properties of Equilibrium

From Eq. (10), we have the following equation for human wealth,

$$H_i(\sigma) = \varepsilon + \sum_{\sigma'} Q(\sigma'|\sigma)\, x_i'\, H_i(\sigma'). \tag{D1}$$

From the definition of total wealth we have that $W_i(\sigma') - H_i(\sigma') = a_i(\sigma')$ where $a_i(\sigma')$ is the amount of Arrow security held by agent i that pays one unit if σ' is realized. Assuming market clearing means that for each σ',

$$\sum_{i=1}^{N} a_i(\sigma') = 0, \qquad \forall\sigma', \tag{D2}$$

and hence, using Eq. (C8), we have that

$$\sum_{i=1}^{N} W_i(\sigma') = N(\sigma')H_i(\sigma') = \beta\frac{1}{Q(\sigma'|\sigma)}\sum_{i=1}^{N}\mathbb{P}_i(\sigma')W_i(\sigma). \tag{D3}$$

Rearranging this equation and factoring $\mathbb{P}_i(\sigma')$ using Eq. (9) gives the following expression for the pricing kernel

$$Q(\sigma'|\sigma) = \beta p(\boldsymbol{x}')\frac{\sum_{i=1}^{N}\mathbb{P}_i(s')x_i'W_i(\sigma)}{N(\sigma')H_i(\sigma')}, \tag{D4}$$

which establishes Eq. (23) from Proposition 2.

Replacing Eq. (D4) in Eq. (D1) and reversing the order of summation gives

$$H(\sigma) = \varepsilon + \sum_{i=1}^{N} W_i(\sigma) \sum_{\sigma'} \left\{ \frac{\beta}{N(\sigma')H(\sigma')} \mathbb{P}_i(s') p(\boldsymbol{x}')\, x_i'\, H(\sigma') \right\}. \tag{D5}$$

Next, cancel $H(\sigma')$ from top and bottom,

$$H(\sigma) = \varepsilon + \beta \sum_{i=1}^{N} W_i(\sigma) \sum_{\boldsymbol{x}'} \left\{ \frac{\beta p(\boldsymbol{x}')\, x_i'}{N(\boldsymbol{x}')} \right\} \sum_{s'} \mathbb{P}_i(s'). \tag{D6}$$

Using the facts that $\mathbb{P}_i(s') = 1$, $\sum_{\boldsymbol{x}'} \left\{ \frac{p(\boldsymbol{x}')\, x_i'}{N(\boldsymbol{x}')} \right\} = 1 - \delta$ and $\sum_{i=1}^{N} W_i(\sigma) = H(\sigma)$ this expression simplifies to,

$$H(\sigma) = \varepsilon + \beta H(\sigma)(1 - \delta), \tag{D7}$$

or

$$H(\sigma) = \frac{\varepsilon}{1 - \beta(1 - \delta)} \tag{D8}$$

which established Eq. (22) in Proposition 2.

Appendix E: Multiplicative Random Process with Reset

Consider the simplest case where, conditioned on survival, returns are IID random variables, i.e.:

$$W_i' = \begin{cases} W_i(1 + \eta) & \text{w.p.} \quad 1 - \delta, \\ 1 & \text{w.p.} \quad \delta \end{cases} \tag{E1}$$

where η is the date t element of a sequence of IID random variables with zero mean and variance equal to σ^2. For this simple case the sequence of conditional probability measures $\varrho(W)$ obeys the operator equation,

$$\varrho(W') = (1 - \delta) \int \mathrm{d}W \varrho(W) \int \mathrm{d}\eta p(\eta)\, \mathbf{d} \left(W' - W(1 + \eta)\right) + \delta \mathbf{d}(W' - 1), \tag{E2}$$

where $\mathbf{d}$ is Dirac's delta function. For large W' this equation delivers a power-law tail, with an exponent μ which is implicitly defined by the self-consistency condition

$$1 = (1 - \delta) \int \mathrm{d}\eta\, p(\eta)\, (1 + \eta)^{\mu}. \tag{E3}$$

In the limit when δ and σ^2 are small, the solution for μ is approximated by the expression,

$$\mu = \frac{1}{2}\left[1 + \sqrt{1 + \frac{8\delta}{\sigma^2}}\right]. \tag{E4}$$

For the wealth process considered in the paper, however, the η are correlated in time (since agent i will consistently make/lose money as long as the sign of $\mathbb{P}_i(t) - \mathbb{P}(t)$ is constant, i.e. during a time $\sim \lambda^{-1}$), and its variance is time dependent (see Eqs. (33) and (34)).

A simplified analysis assumes that η is constant during a time λ^{-1}. This provides the following approximation for μ in this case:

$$\mu \approx \frac{1}{2}\left[1 + \sqrt{1 + \frac{8\delta\lambda}{\bar{\sigma}^2}}\right], \qquad \bar{\sigma}^2 := \mathbb{E}[\sigma^2(t)]. \tag{E5}$$

Note that $\mu \geq 1$ from this formula, meaning that the wealth distribution always has a finite mean when $\delta > 0$.

A way to decrease wealth inequalities is to introduce a wealth tax. If at each time step a small fraction φ of the wealth of each individual is levied and redistributed across the economy, the value of μ in the simple IID model above changes to:

$$\mu = \frac{\varphi + \sqrt{\varphi^2 + 2\delta\sigma^2}}{\sigma^2}.$$

Hence, as expected, increasing φ increases μ and decreases both the Gini coefficient, thereby making markets more efficient in the sense that the difference between $\mathbb{P}$ and $\mathbb{P}_{\text{imp}}$ is reduced.

Appendix F: Proof of Proposition 3

We now seek an expression for the price of a security that pays a dividend d every time $s_t = \{1\}$. This is given by the expression,

$$p_{\text{E}}(\sigma) = \sum_{\sigma'} Q(\sigma'|\sigma)\left[d\,\delta_{s',1} + p'_{\text{E}}(\sigma')\right] \tag{F1}$$

where $\sigma' = (\boldsymbol{x}', s')$ is tomorrow's state, with $\boldsymbol{x}'$ encoding who survives and who dies and $\delta_{s',1}$ is the index function which equals 1 when $s' = 1$ and 0 otherwise. Iterating Eq. (F1) gives the following infinite series:

$$p_{\text{E}}(\sigma) = d\sum_{\sigma'} Q(\sigma'|\sigma)\delta_{s',1} + d\sum_{\sigma',\sigma''} Q(\sigma'|\sigma)Q(\sigma''|\sigma')\delta_{s'',1} + \cdots, \tag{F2}$$

where, from Eq. (29),

$$Q(\sigma'|\sigma) = \beta p(\boldsymbol{x}') \left(\frac{\sum_{i=1}^{N} \mathbb{P}_i(s') W_i(\sigma) x_i'}{N(\sigma')H} \right). \tag{F3}$$

As we have shown in the main text, this object converges, for large N, to

$$Q(\sigma'|\sigma) = \beta p(\boldsymbol{x}') \mathbb{P}_{\text{imp}}(s'), \tag{F4}$$

where

$$\mathbb{P}_{\text{imp}}(s') := \frac{1}{NH} \sum_{i=1}^{N} \mathbb{P}_i(s') W_i(s).$$

Hence,

$$\sum_{\sigma'} Q(\sigma'|\sigma) \delta_{s',1} \equiv \beta \mathbb{P}_{\text{imp}} \tag{F5}$$

where recall that dropping the argument s implicitly means $s = \{1\}$. The first contribution to p_{E} is thus simply

$$d\beta \mathbb{P}_{\text{imp}}.$$

Now let us turn to the second term, which takes the form

$$\sum_{\sigma'} Q(\sigma''|\sigma') Q(\sigma'|\sigma) = \frac{\beta p(\boldsymbol{x}'')}{N(\sigma'')H} \sum_{\sigma'} \sum_{j} x_j(\boldsymbol{x}'') \mathbb{P}_j'(s''|s') W_j'(s')\, Q(\sigma'|\sigma). \tag{F6}$$

Expressing $W_j'(s')$ thanks to Eq. (32), the right-hand side reads:

$$\frac{\beta}{N(\sigma'')H} \left[\sum_{j,\sigma'} \beta x_j(\boldsymbol{x}'') p(\boldsymbol{x}'') \mathbb{P}_j'(s''|s') x_j(\boldsymbol{x}') \mathbb{P}_j(s') W_j(s) + \sum_{j,s'} x_j(\boldsymbol{x}'') p(\boldsymbol{x}'') \mathbb{P}_j'(s''|s') (1 - x_j(\boldsymbol{x}')) H Q(\sigma'|\sigma) \right], \tag{F7}$$

where the first term corresponds to surviving agents in the next time step, and the second term to dying agents that are replaced with new born agents with wealth H.

Consider the two terms of Eq. (F7) in turn. The first term contains a factor

$x_j(\boldsymbol{x}'')x_j(\boldsymbol{x}')$ which equals 1 if an agent j survives for both of the next two periods and zero otherwise. We now use the update rule of agents' beliefs to compute $\mathbb{P}'_j(\sigma''|\sigma')$. One finds, for $s'' = \{1\}$,

$$\mathbb{P}'_j(1|1) = (1-\lambda)\mathbb{P}_j + \lambda; \qquad \mathbb{P}'_j(1|0) = (1-\lambda)\mathbb{P}_j,$$

where we recall that $\mathbb{P}_j := \mathbb{P}_j(1)$. Hence

$$\sum_{s'} \mathbb{P}'_j(1|s')\mathbb{P}_j(s') = [(1-\lambda)\mathbb{P}_j + \lambda]\,\mathbb{P}_j + [(1-\lambda)\mathbb{P}_j]\,(1-\mathbb{P}_j) = \mathbb{P}_j.$$

In words, conditional on survival, the agent's belief is a martingale. Conditioning on $s'' = \{1\}$, one has:

$$\sum_{\boldsymbol{x}'',s''=\{1\}} \beta p(\boldsymbol{x}'') \sum_{j,\sigma'} x_j(\boldsymbol{x}'')\mathbb{P}'_j(s''|s')x_j(\boldsymbol{x}')p(\boldsymbol{x}')\mathbb{P}_j(s')W_j(\sigma) = NH\beta(1-\delta)^2\mathbb{P}_{\text{imp}}.$$

In the large N limit, $N(\sigma'') = N(1-\delta)$ and this term gives a contribution to $p_{\text{E}}(\sigma)$ equal to

$$d\beta^2(1-\delta)\mathbb{P}_{\text{imp}}.$$

Let us now look at the second term. Because of the $1 - x_j(\boldsymbol{x}')$ term, we are looking at states of the world where agent j has died and is replaced by a new agent with an idiosyncratic probability of the next state $\mathbb{P}'_j(s'' = \{1\})$ equal to z, which is uniformly distributed between 0 and 1, with no memory of the past. Therefore, the sum over σ' can be taken independently of the future and gives:

$$\sum_{\boldsymbol{x}'',s''=\{1\}} p(\boldsymbol{x}')x_j(\boldsymbol{x}'')\mathbb{P}'_j(s'' = \{1\}) \sum_{\boldsymbol{x}',s'} (1 - x_j(\boldsymbol{x}'))Q(\sigma'|\sigma) = \beta\delta(1-\delta)\mathbb{E}[z].$$

Hence, we find that dying agents give a contribution to $p_{\text{E}}(\sigma)$ equal to

$$d\beta^2\delta\frac{1}{2},$$

where we have replaced $\mathbb{E}[z]$ by 1/2, and again used the fact that $N(\sigma'') \approx N(1-\delta)$ when $N \gg 1$.

Generalising to all $\ell \geq 1$ time steps in the future, each agent j can either survive ℓ times, with probability $(1-\delta)^\ell$ or die at least once, with probability $1-(1-\delta)^\ell$. In the first case, his/her belief is a martingale. In the second case, the last death cuts all dependence from the past. The calculation above can thus be generalised to give a contribution to $p_{\text{E}}(\sigma)$ equal to:

$$d\,\beta^\ell \left[(1-\delta)^{\ell-1}\mathbb{P}_{\text{imp}} + (1-(1-\delta)^{\ell-1})\frac{1}{2}\right].$$

Summing over ℓ yields our final result for the price of equity in our economy:

$$p_{\mathrm{E}} = \frac{d\beta}{2}\left[\frac{2\mathbb{P}_{\mathrm{imp}} - 1}{1 - \beta(1 - \delta)} + \frac{1}{1 - \beta}\right]. \tag{F8}$$

If agents never die, we recover

$$p_{\mathrm{E}} = d\frac{\beta\mathbb{P}_{\mathrm{imp}}}{1 - \beta},$$

as expected. If agent die at every time step, then $\mathbb{P}_{\mathrm{imp}} \equiv \frac{1}{2}$ and one also recovers the expected result.