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, Revised October 2021

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 and to participants in the University of Virginia macroeconomics seminar in September 2021. 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, Revised October 2021
JEL No. E0,G0,G12,G14

**ABSTRACT**

We construct a model of an exchange economy in which agents trade assets contingent on an observable signal, the probability of which depends on 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. Interestingly, these disagreements generate large wealth inequalities that arise from the multiplicative nature of wealth dynamics which makes successful bold bets highly profitable. 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. The flip side is that such wealth inequalities lead to persistent discrepancies between market implied probabilities and true probabilities.

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

Jean-Philippe Bouchaud
CFM
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 of an exchange economy in which agents trade assets contingent on an observable signal, the probability of which depends on 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. Interestingly, these disagreements generate large wealth inequalities that arise from the multiplicative nature of wealth dynamics which makes successful bold bets highly profitable. 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. The flip side is that such wealth inequalities lead to persistent discrepancies between market implied probabilities and true probabilities.*

> In standard macroeconomic models rational expectations can emerge in the long run, provided the agents' environment remains stationary for a sufficiently long period. [Evans and Honkapohja (2013)].

## I. Introduction

Our opening quote from Evans and Honkapohja encapsulates a commonly held view of macroeconomists: that rational expectations is a justifiable assumption because, in a stationary environment, smart agents are able to learn, after a *sufficiently long time*, about the probability distributions of the economic variables they care about.

A stochastic process is a sequence of random variables; it is stationary if the unconditional probability of an element of the sequence is independent of the date at which it is observed and it is ergodic if averages across possible realisations

* Bouchaud: Capital Fund Management, Chair of Econophysics & Complex Systems, Ecole polytechnique, and Académie des Sciences, 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 and to participants in the University of Virginia macroeconomics seminar in September 2021. 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.

in a given period are equal to the time series average of that variable over many different periods. When ergodicity holds, agents can reliably predict the future by averaging across events that have occurred in the past. Almost all stochastic macroeconomic models are assumed to be ergodic and, for this reason, the argument summarised in our opening quote has proven persuasive to economists who have almost universally adopted the rational expectations assumption since it was introduced into macroeconomics by Robert Lucas fifty years ago (Lucas Jr., 1972).[1]

In the real world, the random events that influence our lives are neither stationary nor ergodic. But although this observation is banal,[2] it is not entirely obvious how to construct a model relevant to economics where ergodicity fails. In this paper, we propose a model to explain why agents fail to learn by exploiting the concept of *quasi-non-ergodicity* widely used in the physics literature to discuss the properties of glasses and "spin glasses" (Anderson, 1989; Debenedetti and Stillinger, 2001). Quasi-non-ergodicity occurs when a stochastic process is ergodic at very long time horizons, but where ergodicity breaks down on a time scale at which realisations from the process might realistically be observed by a human agent.[3] Although the probability distribution of observable variables is ergodic in the long run; as Keynes famously quipped, "in the long-run we are all dead".

Our model differs from previous work in the learning literature in two important respects. First, the stochastic process that agents learn about is self-referential, endowing the process with special properties. Second, we assume that the population changes over time as some agents die and are replaced by newborns with random priors.

To model a self-referential learning process we assume that agents learn the distribution of a public signal which we refer to as *public opinion*. Public opinion is generated as the average probability over the subjective priors of all living agents. To ensure that no agent can exploit ergodicity by being sufficiently patient, we assume that new-born agents do not use the previous history of the public signal. Instead, they begin by making naive forecasts that become increasingly more sophisticated as agents accumulate observations on the signal over time.

We show, in this environment, that it is *reasonable* for agents to infer the probability of the public signal using a common constant gain learning rule with gain parameter $\lambda$. We elaborate on this idea in Section IV.D where we show that,

[1]An alternative related justification for assuming rational expectations dates to the work of Alchian (1950) and Friedman (1953). They argue that, in the long run, only smart actors whose beliefs are closest to the truth will survive. Our results below will contradict that story.

[2]This statement will be obvious to a historian, even if it may not be immediately obvious to the graduate of a modern PhD programme in economics. Although some classes of events may bear superficial similarities to previous events in the same class, no two historical events are identical. History does not repeat itself.

[3]For physical processes such as glasses and spin glasses, the ergodic time scale can be astronomically long at low temperatures. For the model we construct in this paper, it is longer than the life of most individual human beings.

if all other agents learn with gain parameter $\lambda$, using the same learning rule as all other agents generates a forecast that has a negligible bias that is difficult or impossible to detect for almost all values of the time varying probability of the public signal.

We endow our probabilistic world with a market that allows agents to trade two securities that are contingent on realisations of the public signal. At each date, agents solve an inter-temporal optimisation program to determine how much of each security they wish to hold and, because agents have different beliefs, they are willing to trade with each other.

Naively, one would expect the market clearing price to reveal the average 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, and not the unweighted average which corresponds, in our model, to the true probability. Some people accidentally benefit from the mismatch between these two probabilities and temporarily earn higher returns from trading in the asset markets. Interestingly, the resulting distribution of wealth is so unequal that the probability implied by market prices is dominated by the wealthiest agents and fails, even asymptotically, to reveal the true probability of the public signal. Because markets fail to aggregate private information correctly, market prices cannot be used by individuals to reveal the truth.

The wealth distribution generated by our model leads to large and empirically plausible wealth inequalities even though all agents receive the same non-stochastic endowment in every period. We show that the wealth distribution in our model has a Pareto tail as a consequence of the multiplicative nature of wealth accumulation and interestingly, we are able to reproduce the empirical value of both the exponent of the Pareto tail and of the Gini coefficient of real world wealth distributions with reasonable calibrated values of our model's parameters.[4]

## II. 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).[5] 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

[4]A Pareto tail refers to the ability of the Pareto distribution to approximate the density of a non-negative random variable for values that are two or more standard deviations above the mean. A distribution with relatively high probability of observing extreme values is said to possess *a fat tail*.

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

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 (Moran, 1958) in the theory of population dynamics, for which results concerning the time taken to converge to the ergodic distribution were recently obtained by Moran et al. (2020a).

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. Bouchaud (2013) has shown that the Random Field Ising model, which has proven useful in physics to understand interactions between particles, can fruitfully be adapted to understand non-market based interactions between human beings. And Moran et al. (2020b) have shown that ergodicity breaking occurs in models of habit formation. Horst (2017) reviews the literature on ergodicity and non-ergodicity in economic models. In contrast to this literature, we focus on a case where ergodicity is not strictly broken but where the time scale over which it applies may be longer than the lifetime of a human agent. Similar situations are encountered in Agent Based Models (Gualdi et al., 2015) or business cycle models with self-reflexive confidence effects (Morelli et al., 2020).

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 illustrates a different facet of non-ergodicity, where agents adapt their beliefs based on an 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 includes Sandroni (2000) and Blume and Easley (2006) who investigate a "market selection hypothesis", where markets favour traders with more accurate beliefs, and Cogley and Sargent (2008, 2009) who study asset prices in an economy with informed and uninformed agents. In Massari (2019), an interesting scenario is presented where the market selection mechanism fails in the sense that lucky traders become more wealthy than smart traders (as in our model) but prices still manage to remain efficient. In our setting, markets do not favour agents with accurate beliefs and prices fail to reveal the true underlying probabilities.

In an important predecessor to our work, Beker and Espino (2011) study a stochastic endowment economy populated by infinitely-lived Bayesian updaters with heterogeneous priors. They prove that, if markets are complete and endowments are stationary i.i.d. random variables, disagreement vanishes asymptotically with probability 1 if at least one agent's prior contains the support of the true distribution.

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 discrete time stochastic extensions by Farmer et al. (2011) and Farmer (2018) of Blanchard's (1985) perpetual youth model and the stochastic continuous time version of that model in Gârleanu and Panageas (2015). 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 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, agents never agree over their lifetime and the wealth distribution is non-trivial, even in the limit. However, our model has strong philosophical affinities with these papers, in particular Adam et al. (2017), insofar as beliefs depend on past realisations of the stochastic process and lead to market failure, where prices do not reflect fundamentals.

Our story is also in the spirit of the "complex game" model of Galla and Farmer (2013), where in certain conditions agents are unable to reach a Pareto equilibrium through constant gain learning, and become trapped forever in chaotic trajectories that never converge.

## III. Notation and Definitions

In this section we introduce notation and we explain the concepts of stationarity, ergodicity and quasi-non-ergodicity.[6] First, some preliminary concepts.

A *stochastic process* is a sequence of random variables $\{x_t\}_{t\in\mathbb{N}}$ and a stochastic process is *stationary* if the joint probability distribution of

$$(x_{t_1}, x_{t_2}, \ldots, x_{t_k}),$$

is the same as the probability distribution of

$$(x_{t_1+T}, x_{t_2+T}, \ldots, x_{t_k+T}),$$

for all $t_1, t_2, \ldots, t_k, T \in \mathbb{Z}$.

Let $\Omega$ be the space of sequences which take values in a measurable space $\Xi$, with $\sigma$-algebra $\mathcal{B}$, and let $\mathcal{F}$ be the product $\sigma$-field. Define a measure $\mathcal{P}$ on $(\Omega, \mathcal{F})$ which describes the evolution of a process $\{x_t\}_{t\in\mathbb{N}}$ over time. Let $p \in [1, \infty)$ and define $\mathcal{L}^p(\Omega, \mathcal{F}, \mathcal{P})$ as the space of equivalence classes

$$[X] := \{Y : X = Y \ \mathcal{P}\text{-almost everywhere}\}$$

[6] The definitions and organisation in this section draw on the Masters Dissertation by Robinson (2015).

of $\mathcal{F}$-measurable functions $X$ such that

$$\mathbb{E}(|X^p|) < \infty.$$

Define the shift operator $\mathcal{T}$ on $\Omega$ by $(\mathcal{T}\omega)(t) = \xi_{t+1}$, where $\omega(t) = \xi_t$. Then stationarity of the process is equivalent to invariance of $\mathcal{P}$ with respect to $\mathcal{T}$; i.e. $\mathcal{P}\mathcal{T} = \mathcal{P}$. In this case, we call $\mathcal{T}$ a *measure-preserving transformation* for $\mathcal{P}$. We also say that $\mathcal{P}$ is an *invariant measure* for $\mathcal{T}$. Further, define the invariant $\sigma$-field by

$$\mathcal{I} = \{A \subseteq \Omega : \mathcal{T}A = A\}.$$

Given these definitions we have the following *mean ergodic theorem.*

THEOREM 1 (Mean Ergodic Theorem): *Let $p \in [1, \infty)$. Then for any $f \in \mathcal{L}^p(\mathcal{P})$, the limit*

$$\lim_{T\to\infty} \frac{f(\omega) + f(\mathcal{T}\omega) + \ldots + f(\mathcal{T}^{T-1}\omega)}{T} = g(\omega)$$

*exists in $\mathcal{L}^p(\mathcal{P})$. Further, the limit $g(\omega)$ is given by the conditional expectation*

$$g(\omega) = \mathbb{E}_{\mathcal{P}}(f|\mathcal{I}).$$

A stochastic process that satisfies the assumption of the mean ergodic theorem is said to be *ergodic for the mean* and when the conditions of the theorem apply to a stochastic process $\{x_t\}_{t\in\mathbb{N}}$, Theorem 1 implies that sufficiently long time series averages of $x_t$ will converge to the mean of its invariant distribution $P_\infty(x)$, defined as the univariate (single time) marginal of the stationary distribution $\mathcal{P}$. Similar concepts can be used to define ergodicity of higher moments and ergodicity of measures.

When a sequence of random variables is drawn from a stationary ergodic stochastic process, it is possible to make informed conditional forecasts of the future evolution of variables of interest by taking time series averages of past data. The mean ergodic theorem guarantees, for example, that a long time series average of a random variable will converge to the mean of its invariant measure. An example is an AR(1) process generated by the recursion

$$x_t = ax_{t-1} + \xi_t, \quad x_0 = 0,$$

where $\xi$ is a mean zero i.i.d. random variable. For $|a| < 1$, this process is stationary and ergodic and the mean ergodic theorem guarantees that

$$\lim_{T\to\infty} \frac{\sum_{j=1}^{T-1} x_j}{T} \to \mathbb{E}_{P_\infty}[x_T] = 0.$$

But although the theorem guarantees eventual convergence, it does not make any statement about how long convergence takes. That observation motivates the following definition.

DEFINITION 1 (Distance from Equilibrium): *Let the initial value of the random variable be $x_0 = x_{t=0} \in \Xi$, corresponding to an initial distribution $P_0(x)$ with unit mass localised on $x_0$. The distribution of $x_t$ at time $t$ is obtained from $P_0$ as $P_t = \mathcal{T}^t P_0$. The similarity between the conditional distribution $P_t$ and the stationary distribution $P_\infty$ can be characterised by a distance $\mathcal{D} \in [0, 1]$ defined as (Boyd et al., 2004)*

$$\mathcal{D}(P_t, P_\infty | x_0) = \max_{\mathcal{S}} \left| \int_{\mathcal{S}} \mathrm{d}P_t - \int_{\mathcal{S}} \mathrm{d}P_\infty \right|,$$

*where $\mathcal{S} \in \mathcal{B}$ is any subset of $\Xi$.*

The argument of the max operator measures the difference in mass that $P_t$ and $P_\infty$ attribute to any subset of the space $\Xi$. The distance between the two measures, represented by $\mathcal{D}$, is small when $P_t$ and $P_\infty$ assign similar probabilities to all possible subsets of $\Xi$.

We now have enough machinery to define our key concept.

DEFINITION 2 ($\epsilon$-$T$ Quasi Non-Ergodicity): *An ergodic stationary stochastic process $\{x_t\}_{t\in\mathbb{N}}$ is $\epsilon$-$T$ quasi-non-ergodic if*

$$\int \mathrm{d}P_\infty(x_0)\, \mathcal{D}(P_{t>T}, P_\infty | x_0)) < \epsilon.$$

In words, for a given confidence level $\epsilon$, there exists an "ergodic" time $T(\epsilon)$ beyond which one can safely assume that the stationary distribution correctly describes the distribution of the random variable $x$.

The ergodic time $T(\epsilon)$ is related to the so-called mixing rate in the context of Markov chains, and can be computed from the eigenstructure of the transition matrix (see e.g. Boyd et al. (2004)). For the simple AR(1) process above, one can show that the ergodic time is such that

$$\lim_{\epsilon\to 0} \frac{T(\epsilon) \log |a|}{\log(\epsilon)} = 1.$$

As expected, $T(\epsilon)$ becomes infinitely large either when $a \to 1$ (such that the process becomes non stationary) or when $\epsilon \to 0$, i.e. when one demands a strict equality between $P_T$ and $P_\infty$.

In what follows, we will be interested in processes for which the ergodic time is larger than the lifespan of a typical observer, a situation we will simply refer to as *quasi-non-ergodicity.*

## IV. A Two-Outcome, Self-Referential Model

We will build up our argument in three stages.

- In stage one (this section), 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 (sections V, VI and VII), we embed our agents in an endowment economy and we allow them to trade Arrow securities contingent on the realisation of the binary random variable.
- In stage three (section VIII), 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 $\{s_t \in \mathbf{S} \equiv \{0,1\}\}_{t=1}^{\infty}$. This is a simple version of a game that Keynes introduced 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,[7]

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

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

We assume that the event $s_t$ is unrelated to the fundamentals of the economy. In the terminology of Cass and Shell (1983) it is a pure sunspot that we refer to as *confidence*. To make this idea concrete, in Section VIII we provide an interpretation of our model in which the public signal represents a common decision on the part of firms to pay dividends in period $t$. In the absence of self-referential effects, the decision of a firm to pay or not to pay a dividend would be independent of the beliefs of agents. In contrast, in our model the decision of firms to pay a dividend depends on the confidence of market participants about the future, as encoded

[7] More generally, one can consider a model where the true probability is a non-linear, sigmoidal function of the average belief: see Appendix A.A2. Many of the results discussed in the bulk of the paper are actually valid in a more general context, though with interesting twists.

in $\mathbb{P}_t$. Further, the differences of subjective probabilities persist indefinitely, and trigger trades between agents in the securities market.

*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 least-squares learners with different priors, $\mathbb{P}$ does converge to a number in $[0, 1]$, but that number is different for every realisation of $\{\mathbb{P}_t\}_{t=1}^{\infty}$. This setup is an economic analogue of the Pólya urn model discussed in Pemantle (2007) and although the example is instructive, it is not very interesting as a theory of why trades take place in asset markets. Everyone's belief eventually converges to 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 allow the set of decision makers to change over time by recognising that people have finite lives. Second, we replace the assumption of least-squares learning with an alternative constant gain learning algorithm in which people discount the far-away past. This is consistent with the assumption that the world changes: agents are continually replaced and information from the distant past is irrelevant to predictions of the *conditional* probability of events in the near future.

To make these ideas precise, we assume that people die with a probability $\delta$ that 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 V we will use these assumptions to generate simple expressions for aggregate asset prices in a market economy.

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 $\delta$. 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 $\boldsymbol{z}_t \in \mathbf{Z} \equiv [0, 1]^N$ with each of element of $\boldsymbol{z}_t$ is an independent draw from a uniform distribution.[8]

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 $\lambda$ where a value of $\lambda$ closer to 1 means that the person puts more weight on recent outcomes. This term is multiplied by

[8]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. In fact, a slightly different, albeit equivalent specification is to assume that people make occasional observation errors, i.e. mistake $s_t$ for $1 - s_t$.

$x_{i,t}$ 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 VII we will use Eq. (2) to simulate data from an economy with a very 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,

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

where we use the symbol $\mathbf{d}(\cdot)$ to represent the Dirac delta function.[9]

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

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

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

$$\mathcal{P}_t = \mathcal{T}^t\, \mathcal{P}_0,$$

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. More generally, "complex systems" are often defined as probabilistic systems for which probabilities are themselves unknown and must be described with probabilities (Parisi,

[9] A more usual notation is $\delta(\cdot)$ for the Dirac delta function. We use $\mathbf{d}(\cdot)$ to avoid confusion with $\delta$ which we reserve for the age-invariant probability of death.

2007; Bouchaud, 2019).

Introducing the change of variable $u = \mathbb{P} - \frac{1}{2}$, 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 hump-shaped for $\delta > \lambda^2$ and U-shaped when $\delta < \lambda^2$.[10] In our simulations we calibrate the model by choosing $\delta = \lambda^2$, a case for which the invariant measure is *uniform* on $[0, 1]$.

We distinguish between the *memory time* and the *ergodic time*. The memory time is the number of periods beyond which the current observation has a negligible weight on an agent's forecast of the future value of $\mathbb{P}_t$. The ergodic time is the number of periods it takes for time series averages of random variables to approach the mean of the invariant measure. The former is of order $\lambda^{-1}$ and the latter is of order $\delta^{-1}$, as shown in Moran et al. (2020a). By setting $\delta = \lambda^2$ we ensure that the memory time, which for our calibration is approximately one year, is much smaller that the ergodic time, which is approximately fifty years. As a consequence of this difference, people continue to disagree forever. In the limit when $\delta \to \infty$ the invariant measure is a mass point at $\mathbb{P} = 1/2$ and the ergodic time converges to zero.

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 across all members of the population. In symbols,

$$\mathbb{D}_{i,t} \equiv \mathbb{P}_{i,t} - \mathbb{P}_t. \tag{6}$$

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

$$\begin{aligned}\mathbb{D}_{i,t+1} = {} & x_{i,t}\left[(1-\lambda)\mathbb{D}_{i,t} + \delta\left((1-\lambda)\mathbb{P}_t + \lambda(s_t - \frac{1}{2})\right)\right] \\ & + (1 - x_{i,t})\left[z_i - (1-\delta)\left((1-\lambda)\mathbb{P}_t + \lambda s_t + \frac{\delta}{2}\right)\right].\end{aligned} \tag{7}$$

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-$\lambda$ limit, its variance is given

[10] When $\delta \to 0$, the distribution of $\mathbb{P}$ becomes highly peaked around 0 and 1. In fact, such long polarisation periods was Kirman's motivation for introducing his ant recruitment model for opinion dynamics (Kirman, 1993). See also Young (2002). Kirman's model is studied in detail in Moran et al. (2020a) who show that the $\epsilon-$*ergodic time* in this model is proportional to the lifetime of agents, i.e. $T(\epsilon) \propto \delta^{-1}$.

by the expression,

$$\mathbb{V}[\mathbb{D}_i] = \left[\frac{\lambda}{2 + (1-\alpha)\lambda}\right] \frac{\alpha(\alpha+2)}{6(2\alpha+1)} + O(\lambda^3), \qquad \alpha = \frac{\delta}{\lambda^2}. \tag{8}$$

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

$$\mathbb{V}[\mathbb{D}_i] \approx \frac{\sqrt{\delta}}{12}, \qquad (\delta \ll 1). \tag{9}$$

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

In our simulations we chose a time interval of one week and we set $\delta = 3.9 \times 10^{-4}$. 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 is of the same order of magnitude as that reported by Daniel Kahneman (2021) in the assessments of experts using common information. As we will show in our simulations, it is 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.

### D. *Can Some Agents Learn Better Than Others?*

In this section we explore the question: Is it reasonable to use constant gain learning with gain parameter $\lambda$, given that everyone else in the economy is using the same forecast mechanism? We use the word *reasonable* because the use of constant gain learning in this environment is clearly not optimal since the individual learning rule, given by Equation (2),

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

is different from the evolution of the true probability, given by Equation (3),

$$\mathbb{P}_{t+1} = (1-\delta)\left[(1-\lambda)\mathbb{P}_t + \lambda s_t\right] + \frac{\delta}{2}. \tag{3}$$

The difference between the individual learning rule and the true evolution of public opinion arises because agents using naive constant gain learning neglect to account for the arrival of new agents at rate $\delta$.

Consider the problem of a single individual, born at date $j$ who observes the sequence $\{s_t\}_{t=j}^{T}$. The optimal Bayesian forecast of $\mathbb{P}_t$ is the solution to a non-linear filtering problem where $s_t$ provides a noisy signal of the hidden state variable $\mathbb{P}_t$.[11] It follows that constant gain learning is not the best that an arbitrary observer with limitless computational power could achieve. But in the real world people do not have limitless computational power and it may be sufficient to use a simpler rule that has a low predictive error.

So how bad is the constant gain learning rule? Consider the following representation of this rule which we refer to as the R-estimator,

$$R_{t+1} = \sum_{j=0}^{t} (1-\lambda)^{t-j} \lambda s_j + (1-\lambda)^{t+1} R_0.$$

The following discussion is based on an approximation that is valid for time periods in the interval $\lambda^{-1} \ll t \ll \delta^{-1}$ which, for our calibration, is between one year and fifty years. On time scales, where $\lambda^{-1} \ll t$ there is enough data to form estimates of $\mathbb{P}_t$ and for time scales where $t \ll \delta^{-1}$, $\mathbb{P}_t$ is approximately constant.

On these time scales the R-estimator is conditionally biased, with a bias given by the expression,

$$\mathbb{E}[R_t - \mathbb{P}_t | \mathbb{P}_0] \approx \frac{\delta}{\lambda} \left( \mathbb{P}_0 - \frac{1}{2} \right), \qquad \mathbb{P}_0 := \mathbb{P}_{t=0},$$

where the expectation is taken with respect to the true conditional time-varying probability that $s_t = 1$. This bias term reflects the fact that an agent who uses constant gain learning neglects the mean-reverting force towards $\mathbb{P} = 1/2$ which is induced by the birth of new agents with priors centred on $\mathbb{P}_{i,t} = 1/2$.[12]

Could this bias be detected by the agent over time periods of order $\lambda^{-1}$; that is, over lengths of time consistent with the learning window? For that to happen, the agent would need to observe a series of binary outcomes that were statistically implausible given her current belief about the value of $\mathbb{P}_t$. But for most of the range of $\mathbb{P}_t$, the variance of the R-estimator is large relative to its bias. The standard deviation of the R-estimator is approximated by the expression,

$$\mathrm{SD}[R_t] \approx \sqrt{\frac{\lambda}{2} \mathbb{P}_0 \left(1 - \mathbb{P}_0\right)}.$$

To detect the bias, the standard deviation of the estimator must be small relative

[11] Although this problem is superficially similar to the problem of forecasting the state in linear state space model, it is complicated by the facts that the variance of $\{\mathbb{P}_t\}$ is time varying and that the shocks to the state equation and the measurement equation are correlated. The optimal Bayesian forecast could be found using non-linear methods such as the particle filter, but applying methods of this kind are costly.

[12] Note that the unconditional mean of $\mathbb{P}_0$ is equal to $1/2$, which implies that $\mathbb{E}[R_T - \mathbb{P}_T | \mathbb{P}_0 \to 0$ as $T \to \infty$. In words, the $T$-step ahead R-estimator is asymptotically unbiased.

to the bias. This condition is represented by the inequality,

$$\sqrt{\frac{\lambda}{2}\mathbb{P}_0\left(1-\mathbb{P}_0\right)} \ll \lambda\left|\mathbb{P}_0-\frac{1}{2}\right|,$$

where we have used the fact that, for our calibration, $\delta = \lambda^2$ to replace $\delta/\lambda$ on the right side of this expression by $\lambda$. These two terms have a component

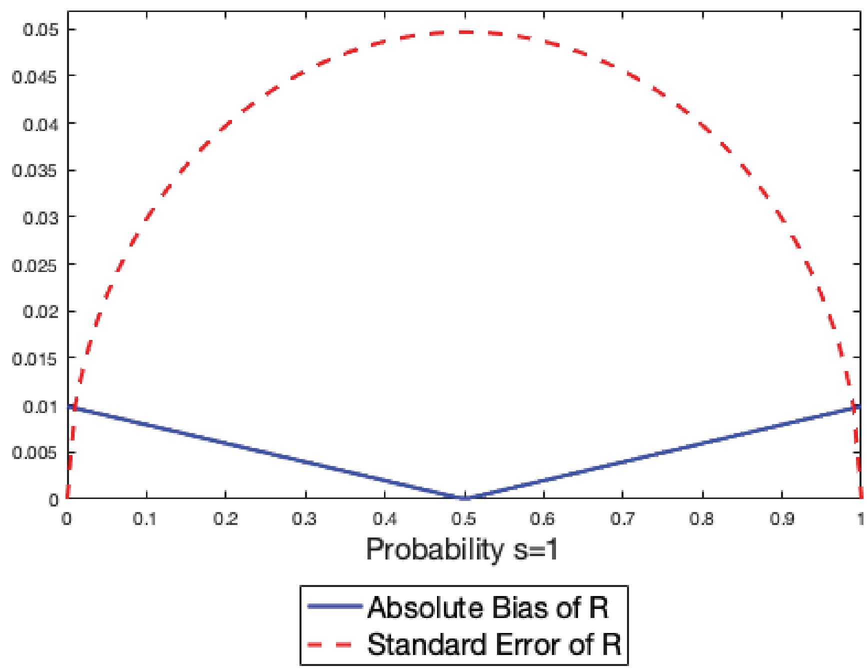


FIGURE 1. A COMPARISON OF THE ABSOLUTE BIAS OF THE R-ESTIMATOR WITH ITS STANDARD ERROR

that depends on $\mathbb{P}_t$ and a component that depends on $\lambda$. In Figure 1 we set $\delta = 3.9 \times 10^{-4}$ and we plot the absolute bias of the R-estimator as the solid line and its standard deviation as the dashed line. Both plots are functions of $\mathbb{P}_t$. This figure makes clear that, except for a sliver of values close to either of the extreme possible values of $\mathbb{P}_t$, the bias of $R_t$ is swamped by its standard deviation.

This result holds because, most of the time, the assumption that $\mathbb{P}_t$ is a random walk is a very good approximation to the truth. But when $\mathbb{P}_t$ gets very close to the boundaries, an observer will begin to observe more mean reverting values than she would consider to be statistically plausible. This anomalous behaviour occurs when $\mathbb{P}_0 < \Delta$ or $\mathbb{P}_0 > 1-\Delta$, where $\Delta$ is a thin sliver of width $\delta^2/2\lambda^3 = O(\lambda)$.

Apart from these rare situations, an observer, using the R-estimator, would not be able to distinguish the small bias in her estimate from measurement noise.[13]

In conclusion, using a simple constant gain learning estimator leads to a negligible bias which is difficult or impossible to detect most of the time. Of course, some smart agents could be aware that the death probability is non zero and account for it in their update rule. However, this would not necessarily make them more successful in the securities market that we will set up in the next section.[14]

## V. Heterogeneous Beliefs in a Market Economy

We have built a model to describe the evolution of public opinion. 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 $\varepsilon$ 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 $\sigma$. We use the adjective *exogenous*, to distinguish the vector $\sigma$ from a vector of *endogenous states* that we introduce in Section V.C. Our agents are irrational in the way they form their beliefs but purely rational in the way they invest and consume.

Rationality and rational expectations are distinct concepts, although they are often conflated in popular accounts of modern macroeconomic models. Our agents are fully rational in the sense that at each point in time they form complete transitive preference orderings over the space of probability measures on future consumption sequences. They *do not*, however, have rational expectations since the probability measures that they use to make choices are subjective. Although these subjective measures do converge to the truth, that convergence occurs at rate $\lambda$ and, for a substantial portion of their lives, agents are trading with incomplete knowledge of the true probabilities of market outcomes.

In the remaining part of this section, we locate our agents in a market economy and we allow them to make trades on all publicly observable events. These events include, not only the binary signal that we refer to as public opinion, but also the realisation of who lives and who dies in every period. This complication introduces $2^N$ new markets since every agent must, in a complete markets economy, trade life insurance contracts contingent on the survival of everyone alive.

Much of Section V involves the introduction of notation to deal with these additional life-insurance markets. Our main results refer to the large $N$ limit and, in this case, there is no aggregate risk from the mortality of individual agents. Although this means that the large $N$ results are much cleaner, we need the finite $N$ machinery to ensure that we account for market clearing constraints in our simulations, where $N$ is finite and equal to a million.

[13]In principle, this bias could be reduced by choosing a slightly larger value of $\lambda$, i.e. a slightly faster rule. But this increases the mean-square error of the estimator. The trade-off between the two would again lead to a negligible improvement of the bias of order $O(\lambda)$.

[14]See the detailed discussion of this point in section VII.D below.

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

$$\mathbb{P}_i(\sigma') = \mathbb{P}_i(s')p(\boldsymbol{x}'). \tag{10}$$

Cass and Shell (1983) distinguish between complete markets, and complete participation in markets. In their sense, our economy has complete markets because every living agent can trade a complete set of securities contingent on all publicly available information. It is, however, an incomplete participation economy because living agents cannot trade with the unborn.

The distinction between complete markets and complete participation is important for the discussion in Section VIII where we show that the equilibrium we construct using Arrow securities can be supported by trades in debt and equity. That result holds only in the large $N$ limit and it relies on the fact that there is no mortality risk in the large $N$ economy.

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

[15]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.

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 IV.B.

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

$$H_i(\sigma) = \varepsilon + \sum_{\sigma'} Q(\sigma'|\sigma)\, x_i'\, H_i(\sigma'). \tag{11}$$

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*

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

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

Eq. (16) is derived by combining equations (11) and (13) with the assumption that agents must remain solvent. $V_i[W_i(\sigma)]$ is the maximum attainable utility given wealth $W_i(\sigma)$, $c_i(\sigma)$ is date $t$ consumption and $\beta$ 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 $H_i(\sigma)$ is defined by the recursion Eq.* (11)*.*

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 $H_i(\sigma')$ 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 by setting 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 3 (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* $y' = \mathcal{G}(y;\cdot)$ *the implied consumption plan solves Problem 1.*

3) *The goods market clears for all* $\sigma'$ *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* $\sigma'$ *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* $H_i$ *is the independent of* $\sigma$ *and is the same for all agents. It is given by the expression,*

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

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.

## VI. 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 VI.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 VI.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[16]

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

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

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

[16]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 $\delta$ remains fixed as $N \to \infty$. Hence $\text{plim}_{N\to\infty} \theta(\boldsymbol{x}') = 1$.

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

$$W_i(\sigma') = x_i' \left[\frac{W_i(\sigma)}{\theta(\boldsymbol{x}')}\right] + (1 - x_i')H. \tag{27}$$

In the large $N$ limit, there is no aggregate mortality risk and, in this case, we obtain the following 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 $\sigma'$ 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.*

Since who dies and who survives is independent from both wealth and beliefs one has, in the large $N$ limit,[17]

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

[17]We use here the fact that if $\eta_i$ and $\xi_i$ 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 $\xi_i = x_i'$.

where we distinguish $N$, which refers to the number of people in state $\sigma$ at date $t$, from $N(\sigma')$, which is the number of survivors in state $\sigma'$ at date $t+1$.

In the large $N$ limit, $\mathbb{P}_{\text{imp}}(\sigma')$ depends on the future realisation of $s$ but not on the mortality state. It is given by the expression,

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

$\mathbb{P}_{\text{imp}}(s')$ is the wealth weighted average probability and it differs from the true probability, $\mathbb{P}(s')$, which is the unweighted 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),

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

$\mathbb{P}_i$ and $W_i$ are *strongly coupled* by the dynamics of individual wealth accumulation, Eq. (32), and because of this strong coupling we cannot split $\mathbb{P}_{\text{imp}}(s')$ into the product of $\mathbb{P}(s')$ and $\text{plim}_{N\to\infty} \sum_i (W_i(\sigma)/N)$ as we did in the common knowledge economy.[18] 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 VIII we will use these expressions to study the implications of our self-referential economy for the prices of debt and equity.

## VII. Results from Simulated Data

In Section VII we illustrate the properties of our model by reporting some statistics for simulated data in a calibrated version. In subsection VII.A we report the results of these simulations and in subsection VII.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 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 2 we graph some data from a single simulation of this calibrated version of our model.

The properties of our simulations are sensitive to two key parameters. The first is $\delta$ which governs the replacement rate of new agents. The second is $\lambda$, which

[18]Strong coupling is a concept borrowed from statistical mechanics where it refers to the failure of a series expansion to converge, even asymptotically.

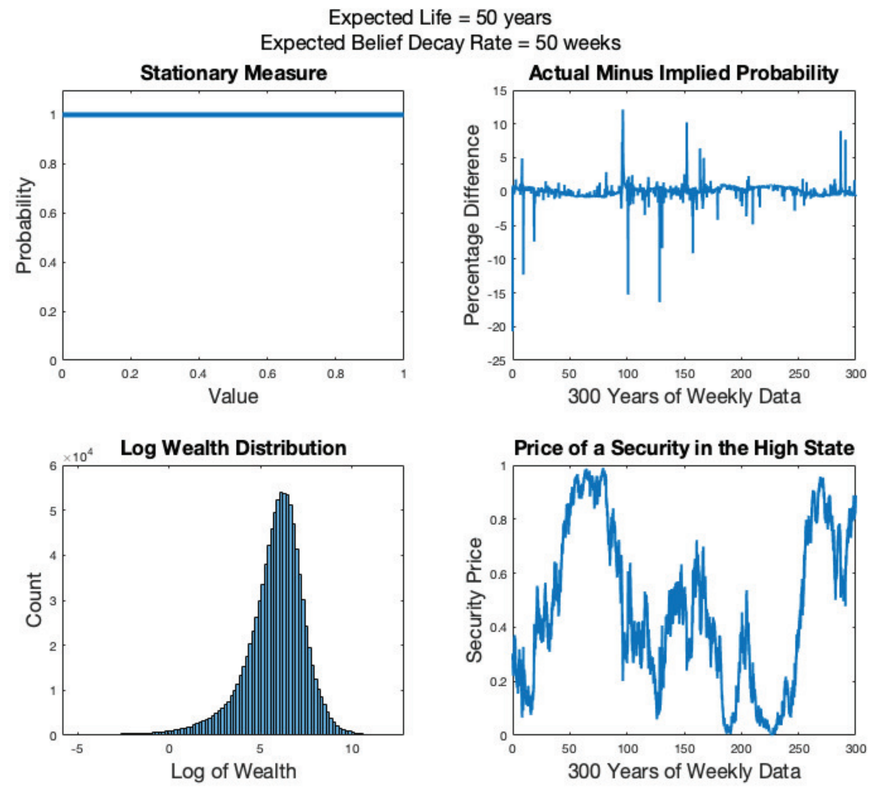


Figure 2. 300 Years of Simulated Weekly Data in an Economy with One Million People

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 $\delta$ 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 $\lambda$ equal to $\delta^{1/2}$, i.e $\alpha = 1$. For our calibrated value of $\delta$, 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 $\delta$ and $\lambda$ 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 $[0, 1]$.

The top left panel of Figure 2 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 $[0, \beta]$ 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 trade by chance, and temporarily, much more successfully 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 VII.C. In Figure 3 we graph the Lorenz curve for the time average of 250 equally spaced samples of the wealth distribution in our simulated data.[19] 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 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 $G(W) \equiv 1 - F(W)$

[19]For large $T$, our sample histogram will converge to the ergodic wealth distribution. There will still be some variability in a sample of 300 years but our experiments with different random draws suggest that this variability not too large.

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

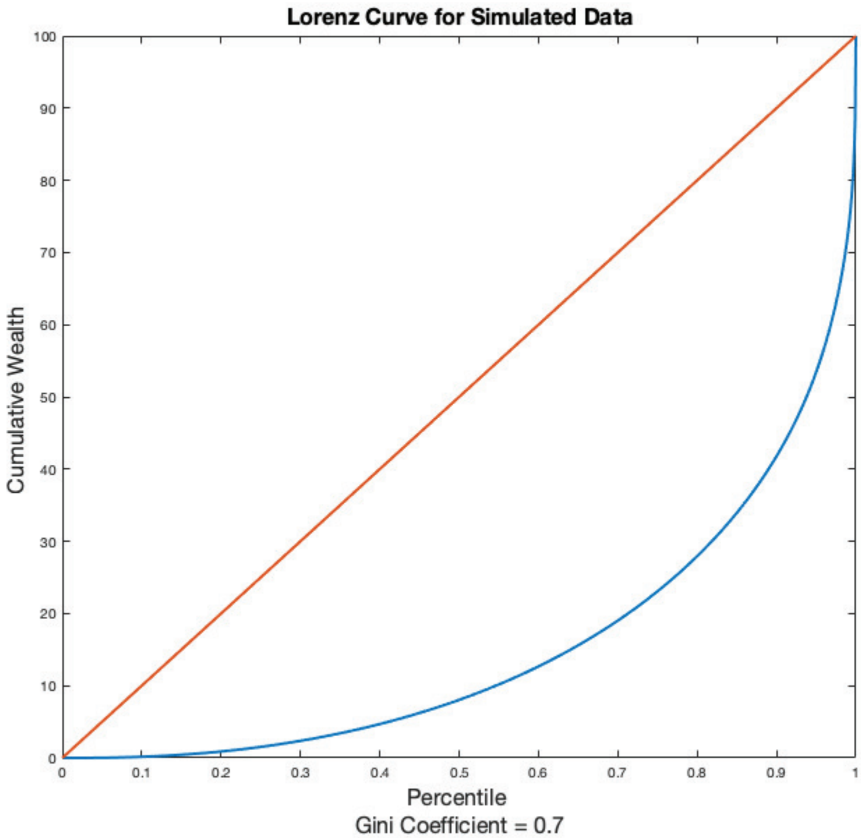


FIGURE 3. THE LORENZ CURVE FOR A SINGLE SIMULATION

| 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[20]

to be the complementary cdf. In Figure 4, 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 $G(W) \sim W^{-\mu}$ corresponds to a probability distribution function (pdf) $\varrho(W) \sim W^{-1-\mu}$. A

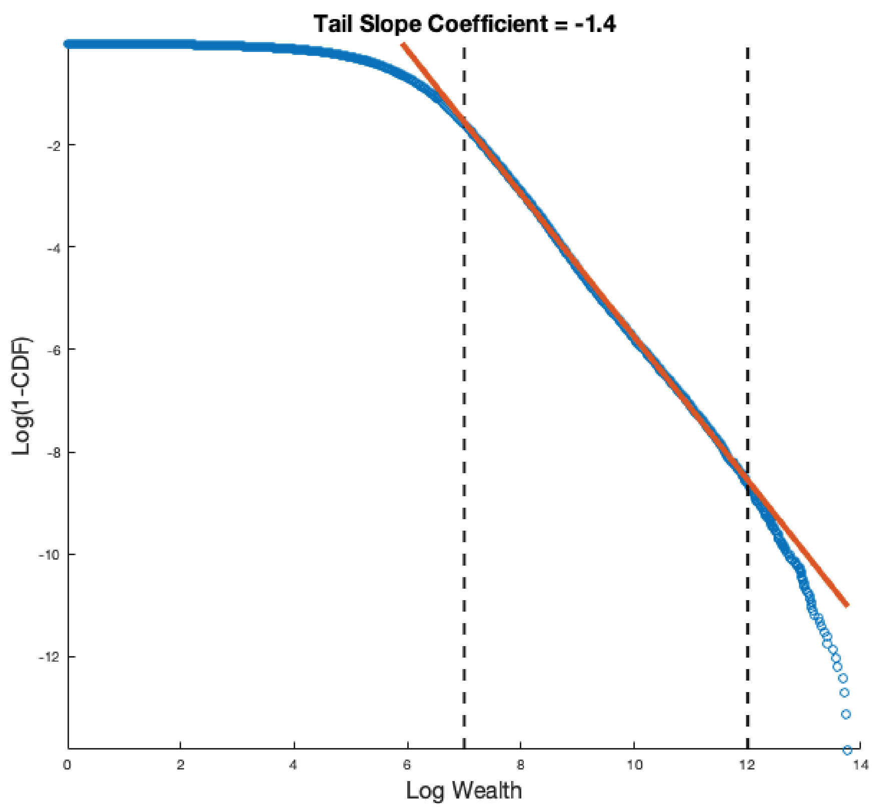


FIGURE 4. ESTIMATE OF THE TAIL PARAMETER IN 1,000 YEARS OF SIMULATED MONTHLY DATA

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.[21] 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

[21]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 $3.9 \times 10^{-4}$. This leads to a value of $H = 1,031$ measured in weeks of income.

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 analysing 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:[22]

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

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 $\mathbb{P}_t$ is also greater than $\mathbb{P}_{\text{imp}}$, and negative otherwise. In fact, provided the sign of $\mathbb{P}_i - \mathbb{P}_{\text{imp}}$ is the same as that of $\mathbb{P} - \mathbb{P}_{\text{imp}}$, the instantaneous expected gain is larger when the bet is bolder, albeit with a larger variance (see Eq. (34) below).

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 – indeed, a most common feature of the real world!

In Eq. (34) we derive an expression for the average of the square of the relative

[22] Eq. (33) follows since

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

change of wealth for surviving agents:[23]

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

One sees from this equation that "bold beliefs", corresponding to a large difference between $\mathbb{P}_i$ and the market probability $\mathbb{P}_{\rm imp}$, 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 $\delta$, 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 $\mu$ depends on the parameters of the problem. We discuss in Appendix E how $\mu$ can be approximately computed. We find, in particular, that $\mu>1$ whenever $\delta>0$.

Random variables with a Pareto tail can be sorted into three classes depending on the value of the tail parameter $\mu$. A Pareto-tailed distribution is well defined for all positive $\mu$ but when $0<\mu\le 1$, the mean and all higher moments do not exist. When $1<\mu\le 2$, 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 $\mu$ 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 belief have a low expected subjective gain from trading. People who disagree may either become spectacularly rich, or spectacularly poor.

[23]Eq. (34) follows from

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

.

### D. The Kelly Criterion

In the introduction, we alluded to the market selection hypothesis which originated with the work of Alchian (1950) and Friedman (1953) and is further developed by Sandroni (2000); Blume and Easley (2006); Beker and Espino (2011) and Massari (2019). The market selection hypothesis is the claim that the agents who survive will be those who maximise the long term growth of their wealth.

When faced with a sequence of gambles, the market selection hypothesis corresponds to the Kelly criterion (Kelly Jr., 1956), which in the current framework amounts to maximising the quantity $\mathbb{E}[\log W_i'/W_i]$. To explore the implications of this hypothesis for our model, let us assume that disagreements are small and define $\upsilon \equiv \mathbb{V}[\mathbb{D}_{i,t}]$, from Eq. (8),

$$\upsilon \equiv \mathbb{V}[\mathbb{D}_i] = \left[\frac{\lambda}{2+(1-\alpha)\lambda}\right] \frac{\alpha(\alpha+2)}{6(2\alpha+1)} + O(\lambda^3). \tag{8}$$

On the assumption that $\upsilon$ is much smaller than unity, we have the following expansion for the log of the growth rate of agent $i$'s wealth,

$$\log\left(\frac{W_i'}{W_i}\right) \approx \left(\frac{W_i'}{W_i} - 1\right) - \frac{1}{2}\left(\frac{W_i'}{W_i} - 1\right)^2 + \dots. \tag{36}$$

Taking expectations of this expression, using equations (33), and (34), one finds the following approximation to the first order in $\upsilon$,

$$\mathbb{E}\left[\log\frac{W_i'}{W_i}\right] \approx \frac{\left[(\mathbb{P}-\mathbb{P}_{\text{imp}})(\mathbb{P}_i-\mathbb{P}_{\text{imp}}) - \frac{1}{2}(\mathbb{P}_i-\mathbb{P}_{\text{imp}})^2\right]}{\mathbb{P}_{\text{imp}}(1-\mathbb{P}_{\text{imp}})} + o(\upsilon). \tag{37}$$

Maximising this quantity over $\mathbb{P}_i$ gives the rule, $\mathbb{P}_i = \mathbb{P}$, which implies that in order to maximise long term growth, the agent should attempt to predict the true probability $\mathbb{P}$, and not the market implied probability $\mathbb{P}_{\text{imp}}$. Making bold bets on the "right" side of $\mathbb{P}_{\text{imp}}$ increases the expected return, but also the variance, which is detrimental to long term growth.

If agents were infinitely lived, the smartest (i.e. the most accurate in terms of predicting $\mathbb{P}$) would become infinitely richer than all other agents and markets would be populated by agents who all agree. In this counterfactual world, $\mathbb{P}$ would converge to $\mathbb{P}_{\text{imp}}$ and agents would never trade. In our model, however, the expected lifetime $\delta^{-1}$ is finite, and the richest agents in any given period are not necessarily the smartest ones but rather those who have made bold, successful bets in the past. It follows that the implied probability $\mathbb{P}_{\text{imp}}$ does not coincide with the true probability $\mathbb{P}$ and agents do not agree with each other, even in the large $N$ limit.

## VIII. Debt and Equity in the Heterogeneous Belief Economy

For much of this paper, we derived results that hold in a complete markets economy by allowing agents to trade a set of Arrow securities. But we motivated our work by referring to the beauty contest quote from Keynes, which refers to trade in long-lived risky assets. In this section we come full circle by demonstrating the applicability of our results to an economy in which agents trade debt and equity.

To make the connection between Arrow securities and the stock market, we assume that $N$ is large and we derive formulae that hold exactly in the large $N$ limit. This assumption allows us to ignore aggregate fluctuations in the annuities markets and to concentrate on trades contingent on disagreement over the realization of $s'$. This signal could be any mechanism for the revelation of public sentiment. What is important for our interpretation is that firms choose to pay dividends only if $s' = 1$.

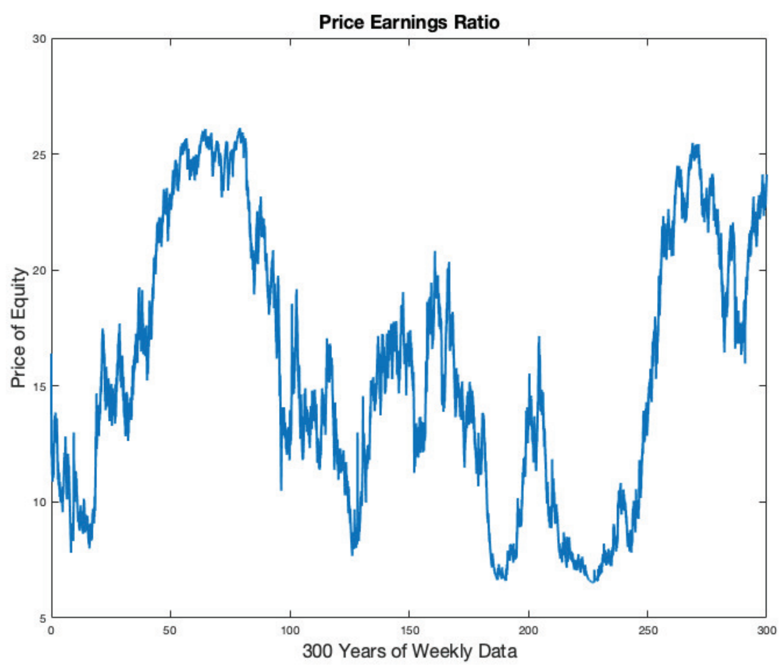


Figure 5. The Behaviour of the Price of Equity in 300 Years of Simulated Monthly Data

We derived explicit trading rules for agents who buy and sell Arrow securities. But there is no reason to restrict ourselves to securities of this kind and the same equilibrium we 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*

$$p_E(s) = \frac{d\beta}{2}\left[\frac{2\mathbb{P}_{imp} - 1}{1 - \beta(1-\delta)} + \frac{1}{1-\beta}\right], \tag{38}$$

$$Q = \beta. \tag{39}$$

For a proof of Proposition 3 see Appendix F. In Figure 5 we have graphed the value of $p_{\mathrm{E}}(s)$ for the data simulated in Figure 2. 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.

## IX. Conclusion

We have a constructed a theory of beliefs in which people exchange information through both market and non-market interactions. Non-market interaction 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 $\mathbb{P}_{\mathrm{imp}}$ but it is unweighted beliefs, $\mathbb{P}$, 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 would be sub-optimal. Quite remarkably, the coupled dynamics of individual wealth and beliefs leads to a fat-tailed distribution of wealth. The richest agents at a given instant in time are not necessarily the smartest ones but rather those who have made bold, successful bets in the past. Since those agents dominate the market, the implied probability $\mathbb{P}_{\mathrm{imp}}$ cannot be used to learn the true probability $\mathbb{P}$.[24]

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

Why are there no Warren Buffets who invest for the very long run by guessing that the probability of a successful outcome will be equal to the mean $\mathbb{P} = 1/2$ of the invariant distribution? Our answer is that this would only be the case if we lived forever and could afford to be strict Bayesian learners, but the world that we live in is far better approximated by observing recent realisations than by relying on an unconditional long-run ergodic measure. We believe that our quasi-non-ergodic model aptly illustrates what Keynes had in mind when he wrote that "In the long run we are all dead".

## 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.

Armen Alchian. Uncertainty, evolution and economic theory. *Journal of Political Economy*, 58:211–221, 1950.

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

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.

Jean-Philippe Bouchaud. Crises and collective socio-economic phenomena: simple models and challenges. *Journal of Statistical Physics*, 151(3-4):567–606, 2013.

Jean-Philippe Bouchaud. Radical complexity. *Royal Economic Society*, 2019. URL `https:// www.res.org.uk/ resources-page/ radical-complexity.html`.

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.

Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004. doi: 10.1137/S0036144503423264. URL https://doi.org/10.1137/S0036144503423264.

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

Dave Cass and Karl Shell. Do sunspots matter? *Journal of Political Economy*, 91:193–227, 1983.

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.

Cass R. Sunstein Daniel Kahneman, Olivier Sibony. *Noise: A Flaw in Human Judgment*. Little, Brown Spark, New York, first edition, 2021.

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.

George W. Evans and Seppo Honkapohja. Learning as a rational foundation for macroeconomics and finance. In Roman Frydman and Edward S. Phelps, editors, *Rethinking Expectations: The Way Forward for Macroeconomics*, pages 68 – 111. Princeton University Press, Princeton, 2013.

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.

Milton Friedman. *Essays in Positive Economics*. University of Chicago Press, Chicago, 1953.

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.

Tobias Galla and J. Doyne Farmer. Complex dynamics in learning complicated games. *Proceedings of the National Academy of Sciences*, 110(4): 1232–1236, 2013. ISSN 0027-8424. doi: 10.1073/pnas.1109672110. URL https://www.pnas.org/content/110/4/1232.

Nicolae Gârleanu and Stavros Panageas. Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy*, 123(3):670–685, 2015. doi: 10.1086/680996. URL https://doi.org/10.1086/680996.

Stanislao Gualdi, Marco Tarzia, Francesco Zamponi, and Jean-Philippe Bouchaud. Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50:29–61, 2015. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2014.08.003. URL https://www.sciencedirect.com/science/article/pii/S0165188914001924. Crises and Complexity.

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.

John Larry Kelly Jr. A new interpretation of information rate. *Bell System Technical Journal*, 35(4):917–926, 1956.

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.

Robert E. Lucas Jr. Expectations and the neutrality of money. *Journal of Economic Theory*, 4:103–124, 1972.

Filippo Massari. Market selection in large economies: A matter of luck. *Theoretical Economics*, 14(2):437–473, 2019. doi: https://doi.org/10.3982/TE2456. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/TE2456.

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.

G. Parisi. Physics, complexity and biology. *Advances in Complex Systems*, 10(supp02):223–232, 2007. doi: 10.1142/S021952590700132X. URL https://doi.org/10.1142/S021952590700132X.

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.

Benjamin Robinson. Ergodicity of stochastic processes and the markov chain central limit theorem, 2015. URL https://people.bath.ac.uk/br275/projects/MastersProject.pdf.

Alvaro Sandroni. Do markets favor agents able to make accurate predictions? *Econometrica*, 68(6):1303–1342, 2000.

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

H. P. Young. *Individual strategy and social structure*. Princeton University Press, 2002.

## 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:

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

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 $\lambda^3$ yields:

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

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 $\lambda^3$, and we will neglect them in the following approximation.

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

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

This equation coincides with the continuous time description of Kirman's ant recruitment model (Kirman, 1993), for which a lot is known (see Moran et al. (2020a) for recent results and references).

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

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

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 $\beta$, $\mathcal{V}(x)$ 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 $\phi$ 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 $\Gamma$ can be fully computed (at least numerically) for *any* potential $\mathcal{V}(x)$. The exponential dependence of $\tau_\times$ in $\lambda$ 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. (7) 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. (7) 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 $\alpha$ 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. (8).

## 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. (11), 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 $\sigma'$ is realized. Assuming market clearing means that for each $\sigma'$,

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

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. (10) 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 $\eta$ is the date $t$ element of a sequence of IID random variables with zero mean and variance equal to $\sigma^2$. 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 $\mu$ 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 $\delta$ and $\sigma^2$ are small, the solution for $\mu$ 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 $\eta$ 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 $\eta$ is constant during a time $\lambda^{-1}$. This provides the following approximation for $\mu$ 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 $\varphi$ of the wealth of each individual is levied and redistributed across the economy, the value of $\mu$ in the simple IID model above changes to:

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

Hence, as expected, increasing $\varphi$ increases $\mu$ 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,

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

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:

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

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 $p_{\text{E}}$ 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 $\sigma'$ 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 $\ell$ 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 $\ell$ 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.