THE UNIVERSITY OF CHICAGO

THE STATISTICAL MECHANICS OF LEARNING AND CO-ADAPTATION IN HETEROGENEOUS POPULATIONS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
JORDAN TIMOTHY KEMP

CHICAGO, ILLINOIS
DECEMBER 2024

To my mother, Tamu, who honed my curiosity and creativity

To my father, Shane, who channeled these into good work

If you surrender to the wind, you can ride it.

-Toni Morrison

TABLE OF CONTENTS

LIST OF FIGURES

*

ACKNOWLEDGMENTS

Thank you to the following educators and scientists who instilled in me the knowledge, the skills, the technique, and most importantly the confidence to undertake a PhD. They are James Brown, Dr. Ed Myers, Katherine Metrick, Prof. Roger Tobin, Prof. Moon Duchin, Prof. Hannes Bernien, and Prof. Kevin Singh.

Thank you to my earnest friends, without whom I would have struggled through my doctoral coursework. They include Akhil Ghanta, Amanda Farah, Adrien Sy, and Adam Kline.

Thank you to my research colleagues from whom I drew endless inspiration, motivation, and enthusiam for research. They include Dr. Shankar Menon, Dr. Andrew Stier, Prof. Marc Berman, Andrew Pocklington, Prof. Olivier Gallay, and Prof. Max-Olivier Hongler.

Thank you to the research and support staff who helped me find home and community at the university, who maintained my focus, and who would not let opportunities for personal and professional development pass me by. They include Putri Kusumo, Dr. Brian Wilson, Dr. Julia Koschinsky, and the jubilant Aimee Giles.

Thank you to my committee, Prof. Sidney Nagel, Prof. Stephanie Palmer, and Prof. Arvind Murugan.

Lastly, thank you to Luís Bettencourt who opened my eyes, who provided opportunity, encouragement, and patience. Who gave me freedom, independence, and trust. Who impressed the power of elegance and simplicity, who inspired healthy skepticism, and effused methodical optimism. It is plain luck to find advisor so compatible in aspirations and interest, but glimpses of the coming road and its anticipated doors reveal how truly fortunate I have been to have his advice for these four years.

ABSTRACT

Adaptation is a fundamental mechanism of growth. Scientists have developed statistical models in numerous contexts to characterize growth and its emergent behaviors, such as inequality, competition, and cooperation. However, we still lack a general adaptive mechanism that explains the emergence of growth in uncertain environments, preventing systematic exploration of the origins of agent heterogeneities. In this dissertation, I derive a theory of statistical growth among agents adapting to their environments. I then show several key results. First, that the average growth rate of agents' resources is governed by the information they hold about their environment. It follows that the learning process can attenuate growth rate disparities, reducing the long-term effects of heterogeneity on inequality. Second, I show how groups that optimally combine complementary information about resources maximize their effective growth rate. I show that these advantages are quantified by the information synergy embedded in the conditional probability of environmental states given agents' signals, such that groups with a greater diversity of signals maximize their collective information. Lastly, using simple, pairwise agent interactions, I show how agent preferences converge when driven by observation of each other's behaviors. These results demonstrate how the formal properties of information underlie the statistical dynamics of many complex processes across biological and social phenomena.

CHAPTER 1

INTRODUCTION

Adaptation refers to the changes in the function or behavior of a system, such as an organism, economic agent, or ecological community, that enhance its ability to survive in specific environmental conditions. Environments vary widely and often promote distinct behaviors. In social environments, built from the interactions of many agents with individual, heterogeneous preferences, agents adapt to reflect social norms, behaviors, or institutions. In natural or professional settings, agents are rewarded for filling specific niches, such as specialized skills in supply chains or a specific biomass conversion in energy cycles. Across these contexts, adaptation is driven by the pursuit of resources like social capital, or material resources such as money or energy. Generally, agents that are better adapted and exhibit more favorable behaviors can survive better by extracting more resources over time.

Social and biological scientists have long sought insight into how this process unfolds across scales, from individual behaviors to collective decisions. Formally, they study how the properties of complex systems, like economies or ecosystems, emerge from the heterogeneous strategies and social structures their interacting constituents adopt in response to their dynamic environments. They are also often interested in whether adaptation produces strategies that are optimal and result in the highest rewards possible given the environment. In general, this is difficult to do in a complete and mathematically rigorous way. However, improvements to data collection methods and computational power in the mid-20th century have revolutionized how we understand and explore these emergent behaviors, inspiring new theoretical approaches. A wealth of models inspired by statistical mechanical systems out of equilibrium has revealed connections between the dynamics of strategic agents and macroscopic properties of population growth, inequality, and cooperation.

These dynamical models typically describe a general agent, a proxy for a decision-maker, such as a bacterium foraging for energy, a rabbit evading predators, or a person investing

their wealth. They describe behaviors over time and are usually nonlinear [28, 85, 204]. Some processes like reproduction can be multiplicative, while others may involve higher-order interactions or feedback processes [6, 12, 81]. This complexity limits the insights that analytical calculations can provide, prompting researchers to employ computational and numerical methods. These approaches typically describe a deterministic process that captures the agent's mean tendencies, such as an interaction preference or the average payoff of an action. This represents what we know or assume about the agent's behavior and can be adjusted to study emergent population-level phenomena over time. A statistically distributed, or stochastic, term is also included to account for unpredictable behaviors that are difficult to model precisely but have predictable statistical patterns. To illustrate these concepts in a formalized manner, consider Brownian motion, a canonical example of a stochastic dynamical process.

Consider a particle with position vector $\boldsymbol{X}$. At every time step dt, the particle drifts on average $\boldsymbol{\gamma} dt$ with some statistical fluctuation $\sigma d\boldsymbol{W}_t$. Here, $d\boldsymbol{W}_t$ describes a Wiener process, or a mean-free Gaussian distributed fluctuation $d\boldsymbol{W}_t \sim \mathcal{N}(0, 1)$. The stochastic differential equation (SDE) is given by

$$d\boldsymbol{X} = \boldsymbol{\gamma} dt + \sigma d\boldsymbol{W}_t.$$

In a physical context, this could represent the drift of a charged particle, where $\boldsymbol{\gamma} \propto q\boldsymbol{E}dt/m$ for some charge q of mass m in an electric field $\boldsymbol{E}$, with fluctuation $\sigma \propto \sqrt{2dk_BT}/m$ for drag d, the Boltzmann constant k_B at temperature T. [1] In this treatment, assume the particle is subject to a measurable force, given by $q\boldsymbol{E}$, and collides with innumerable other particles. These collisions result in a diffusion process that scales with the electron mobility d and as a function of temperature T such that electrons that are hotter and move faster in response to an electric field diffuse further. Although we cannot track individual collisions,

1. This mapping is far from exact, and is used only for analogy. The actual Langevin equations model the position and velocity as coupled SDEs as the charged particle velocity is also stochastic, whereas here it is not. The full description is provided in [224], where the Langevin definition is given in equation 6.3.

we can describe the statistical properties of the particle's trajectory. Usually, this noise is diffusive and drives the system towards disorder, but it may also provide insights into how these stochastic dynamics behave. For example, fluctuations could have a particular functional form representing the noise generating process. Noise could be correlated across space, if dynamics between agents are coupled, or time, if past states affect future behaviors. This model will serve as a starting point for the theory explored in Chapter 2, and will be expanded upon later in the text. However, this dissertation will only explore cases where noise is independent and identically distributed (i.i.d).

Nonequilibrium models with analogs to Brownian Motion are often used to study adaptive dynamical systems, because they capture both random fluctuations and deterministic interactions. For instance, the stochastic Lotka-Volterra equations use coupled multiplicative Brownian populations to examine in which parameter regimes predator-prey dependencies lead to oscillating or stable dynamics ([161, 226]. Wealth inequality is often modeled as exponential Brownian motion with mean-field interaction terms representing wealth redistribution or transfer ([39]). Meanwhile, networks of beliefs over coupled, binary opinions rely on more complex models, such as Ising lattices ([96]).

While these models have successfully identified and characterized several common emergent phenomena, we lack theory on how individual behaviors and adaptations influence models of complex systems like predator-prey dynamics or functions of the brain. A critical piece missing from these models is a description of how these adaptive behaviors develop in the context of their environment. That is, these system emerge from adaptive processes (i.e., the need to manage resources to survive), but the models often studied do not specify specific adaptive mechanisms. Decoupling agent behaviors from this fundamental drive towards optimality limits modelers to *ad-hoc*, phenomenological explanations for behavioral mechanisms (e.g. *where do mean and noisy behaviors come from?*), and blinds modelers from fundamental considerations (e.g. *how do environments influence behavior?*). More importantly, this

decoupling prevents modelers from leveraging the deep connection between information and behavior ([152, 204]), and the rich theoretical insights that an information theory provides [89, 184]. This raises the question: how does adaptation in a complex environment determine the dynamics of growth and social behavior?

1.1 Dissertation overview

This dissertation is organized as follows. Chapter 2 motivates the need for an adaptive theory of growth by demonstrating the shortcomings of standard Geometric Brownian Motion (GBM) approaches to studying wealth dynamics. I show that parameter heterogeneity determines the properties of aggregate growth and drives inequality over long timescales. With the need for a better theory established, Chapter 3 examines the origins of heterogeneity in terms of agent decision-making. It derives statistical mechanics for agents' resource growth and inequality in noisy environments, grounded in Shannon information, further advancing the biophysical and portfolio theory literature. The chapter concludes by exploring the roles of environments and adaptability in promoting and mitigating inequality. To illustrate the insights introduced by an information-theoretic approach, the remaining two chapters build on these dynamics to derive interactions in cooperative and competitive contexts. Chapter 4 derives the benefits of skill complementarity across coordinating agents to growth in terms of information synergy. This result shows how information theory explains the promotion of cooperation in heterogeneous populations, introducing a novel cooperation mechanism to the literature. Chapter 5 then provides a primer for studying competitive dynamics. It derives the convergence properties of preferences of agents adapting to each other's behavior through Bayesian inference based on observations of past behaviors. Chapter 6 concludes this dissertation by summarizing its contributions. It then proposes three areas of inquiry for future research that will enhance our ability to study adaptive dynamics: adaptation in environments with more decision-making options and varied payoffs, multimodal coopera-

tion across different group sizes, and competition dynamics in larger groups with various environmental interactions.

CHAPTER 2

DYNAMICS OF GROWTH AND INEQUALITY IN ADAPTIVE POPULATIONS

2.1 Introduction

Exponential growth and broad inequality are general features of adaptive population dynamics in biology and society [213]. In biology, the maintenance of diverse populations is associated with greater biodiversity and with larger pools of variability enabling faster processes of evolution by natural selection [88]. In human societies, rates of long-term (economic) growth are much higher than in most ecosystems, often generating widening inequalities and leading to familiar stresses of social justice and equity [62, 197] [1].

Researchers are particularly interested in the dynamics of economic inequality [213], applying multidisciplinary approaches that include historical data analyses [135] and searches for social, political, and economic mechanisms that generate, distribute, and rebalance wealth within populations [71, 197]. Research in economics has emphasized the importance of heterogeneity in populations [62] as a source of widening wealth inequality. These heterogeneities are associated with a number of different population features such as the divergence of incomes between capital and labor [197], between management and workers within firms [7, 9], and between people with different educational attainment [102]. These primarily empirical approaches point to the fundamental importance of diverging income growth for different subpopulations and to the need for rebalancing mechanisms if inequality is to remain controlled. However, we still do not have statistical theories of wealth dynamics applicable in general circumstances. As we will see by the end of this chapter, we lack a clear picture of how hetergeneous growth emerges from behavior or experiences, and therefore lack insight

1. This chapter is adapted from Kemp, J. T., & Bettencourt, L. M. (2022). Statistical dynamics of wealth inequality in stochastic models of growth. *Physica A: Statistical Mechanics and its Applications, 607,* 128180.

on how to concurrently manage growth and inequality.

To this end, the application of nonequilibrium statistical physics models to the study of wealth dynamics has produced a number of results on the role of stochasticity as a source of inequality. These models have explored inter-agent exchanges (particularly in the mean-field limit), and redistribution mechanisms in the context of stochastic geometric growth models [38, 47, 67, 97, 120, 170, 188, 225]. The main focus of this work has been the design of wealth redistribution schemes towards creating long-time stationary limits [21, 39, 159] with parametrically controlled levels of inequality and whether such stationary solutions exist [22, 195].

While these models have been used to fit data on wealth distributions [22, 68, 159, 209], their specialized approaches have left open questions on the fundamental statistical dynamics underpinning the generation and allocation of wealth. For example, evidence suggests models following strict application of Gibrat's law (individual growth rates independent of wealth) cannot characterize the rapid emergence of inequality experienced in recent years [95, 97], motivating explicit analysis of open-ended dynamical effects due to various sources of fluctuations and correlations. These developments support multiplicative stochastic growth as a starting point for modeling [23], but call for the consideration of statistical effects due to fluctuations and correlations in model parameters, especially growth rates.

To address these issues, this chapter derives the time evolution of the distribution of resources, a proxy for wealth, in a statistical population of heterogeneous agents experiencing geometric random growth. It focuses on the dynamic effects of growth rate fluctuations and their correlations with resources across agents and over time. I show that these two effects lead to two different dynamical time scales, which require different control measures so that inequality does not explode in a population over time. Specifically, I show that natural schemes to reduce inequality are subject to reversal over longer time scales because of the variability of growth rates in the population. I end this chapter by discussing the

mechanisms that may simultaneously lead to sustained exponential growth and control of long-term sources of inequality and propose redistributing growth opportunities through agent-based adaptative decision making in stochastic environments, which will be detailed in Chapter 3.

2.2 Geometric Brownian Motion in Correlated Populations

The fundamental model for the dynamics of resources in populations relies on stochastic exponential growth [23, 95]. In its simplest form, known as geometric Brownian motion, agents generate wealth by (re)investing incomes net of costs. Crucially these models, unlike additive stochastic growth, generate population lognormal and power-law statistics, which characterize observed wealth and income statistics [53, 210]

Specifically, the dynamics of stochastic multiplicative growth start by tracking the resources (wealth), $r(t)$, of a specific agent at time t. This quantity changes in time via the difference between an income, $y(t)$, and costs $c(t)$. The difference, $y-c$ (net income), is then defined to be proportional to resources, expressed as a (potentially r dependent) stochastic growth rate, $\eta(t) = (y(t) - c(t))/r(t)$, for $r > 0$. It follows that the time evolution of resources obeys a simple Gaussian multiplicative stochastic process: $\frac{dr}{dt} = \eta(t)r$, where $\eta(t)$ is a stochastic growth rate with temporal mean $\bar{\eta}$, and temporal variance σ^2, where the standard deviation, σ, is known as the volatility. When these two time averages are independent of r and t, this stochastic equation can be integrated via Itō calculus [23] to give

$$\ln \frac{r(t)}{r(0)} = \left(\bar{\eta} - \frac{\sigma^2}{2}\right)t + \sigma W(t), \tag{2.1}$$

where $\gamma \equiv \bar{\eta} - \sigma^2/2$ is the effective growth rate and r_0 the agent's resources at $t = 0$, $r_0 = r(0)$. The decrease in the mean growth rate $\bar{\eta}$ due to finite volatility is an important feature of multiplicative growth. The quantity $\sigma W(t)$ is a Wiener process with magnitude

proportional to the volatility and units of $t^{-1/2}$.

2.2.1 Average Growth Rate in Correlated Populations

Under these circumstances, single-agent dynamics produce the population dynamics of a noninteracting community of agents with uniform growth rates and volatilities. That is, for homogeneous wealth value r and growth rate $\bar{\eta}$, the sample average growth over a finite population of size N is given by $\langle \bar{\eta} r \rangle_N \equiv \frac{1}{N} \sum_i \bar{\eta}_i r_i = \bar{\eta} r$. This basic case is unrealistic as wealth and income vary across real populations. To deal with this more general situation, I introduce heterogeneity to the population by sampling agents from distributions of $\bar{\eta}$, $\langle \bar{\eta} \rangle_N = G$, and r, $\langle r_0 \rangle_N = \mu_0$. When the parameters are uncorrelated and sampled from independent distributions, the population average growth is the growth averaged over the population, $\langle \bar{\eta} r \rangle_N = G\mu_0$. However, when initial wealth r and growth rates $\bar{\eta}$ are correlated as historical wealth distribution analysis would suggest [135], growth of the average separates into the population mean term and a covariance term given by (A.1)

$$\langle \bar{\eta} r \rangle_N = G\mu_0 + \mathrm{cov}_N(\bar{\eta}_i, r_i) \tag{2.2}$$

for parameter covariance cov_N.

The growth term of the correlated expectation value, $G' \equiv G + \mathrm{cov}_N\big(\bar{\eta}_i, \frac{r_i}{\mu_0}\big)$ is called the population averaged resource growth rate. This quantity can be calculated analytically for Gaussian distributed growth rates and log resources, $\ln r_0$, with population variances $\langle (\bar{\eta} - G)^2 \rangle_N = \sigma_{\bar{\eta}}^2$ and $\langle (\ln r_0 - \ln \mu_0)^2 \rangle_N = \sigma_r^2$. The result involves the variances of both quantities as well as the Pearson correlation coefficient between them, as (A.1)

$$\mathrm{cov}_N\left(\bar{\eta}_i, \frac{r_i}{\mu_0}\right) = \rho \sigma_{\bar{\eta}} \sigma_r \exp\Big[\sigma_r^2/2\Big]. \tag{2.3}$$

The transformation to lognormal distributed resources is performed for convenience in later

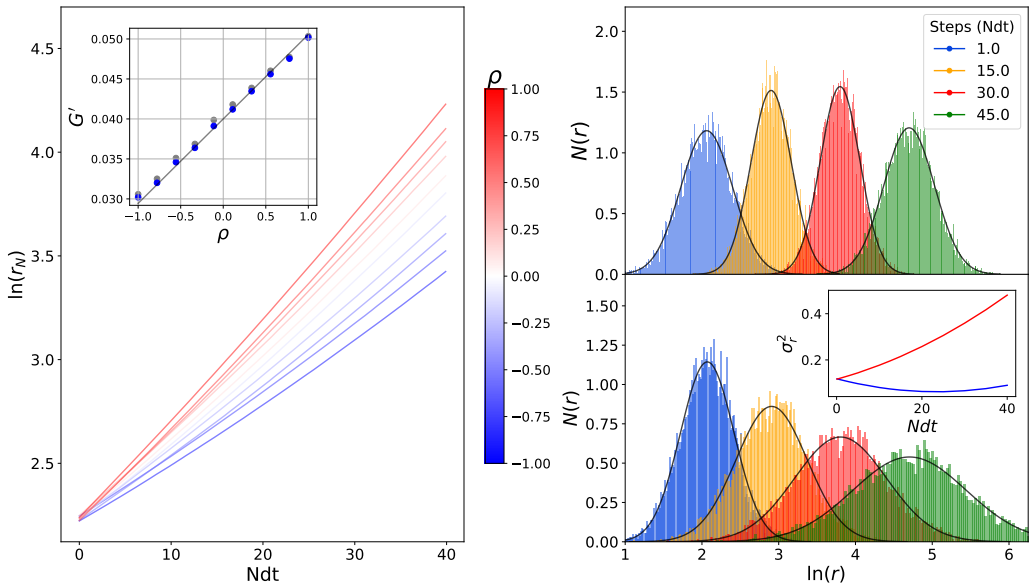

Figure 2.1: Monte Carlo simulations of growth dynamics with correlated growth rates and initial resources. *Left:* Aggregate resource trajectories of populations with ρ ranging from -1 (blue) to 1 (red). *Inset*: The population-averaged effective growth rate (blue) scales linearly in ρ, following Eq. 2.4. The temporal averaged growth rate (gray) scales similarly. *Right:* Distributions of log wealth at different time steps for a regressive, $\rho = .75$ (bottom), and progressive, $\rho = -.75$ (top) population. Black curves represent normal distribution fits. The regressive assignments broaden more quickly, indicating that the populations are becoming more rapidly unequal. *Inset:* Variance in log resources of the progressive (blue) and regressive (red) populations over many time steps. Data are simulated from populations of 10^5 agents with Gaussian distributed initial resources of mean $\ln \mu_0 = 2$, mean growth rate $\bar{\eta} = .04$, and standard deviations $\sigma_{\bar{\eta}} = .015$, $\sigma_r = .682$.

calculations, and introduces the exponential term to the resources in the covariance (if instead resources were normally distributed, the exponential term becomes unity). Thus, the population average effective resource growth rate, Γ', is expressed in terms of the model parameters as

$$\Gamma' = G - \frac{\sigma^2}{2} + \rho\sigma_{\bar{\eta}}\sigma_r \exp\left[\sigma_r^2/2\right], \tag{2.4}$$

where Γ is reserved for the mean effective growth rate $\Gamma \equiv G - \sigma^2/2$.

This expression shows that a positive covariance between initial log resources and growth

rates results in a higher population averaged effective growth rate and vice versa. Naturally, associating higher-growth opportunities with wealthier individuals will exacerbate inequalities in the population over time, whereas the opposite would reduce inequality, at least over short time scales. As such, I denote population dynamics with positive correlation as *regressive*, and with negative correlation as *progressive*. I visualize the tradeoff between short-term change in inequality and growth through numerical simulations in Fig. 2.1.

Conditions for Progressive Population Dynamics

We see that these simple considerations present an apparent paradox for any attempt to simultaneously maximize total wealth, a measure of average social welfare, and reduce inequality. Decreasing ρ, thereby making the societal dynamics more *progressive* results in a social opportunity cost in terms of a decline in average growth. This effect may even lead to negative exponential growth, and thus to a decline in societies starting out with low average growth and high volatility. One way out of this dilemma is for a progressive policy assignment to create, in ways to be specified, higher average growth rates than the regressive case.

To analyze this possibility, Eq. 2.4 introduces the threshold correlation below which the average effective resource growth rate becomes negative. Computed under the condition $\Gamma' = 0$, this critical correlation value (denoted with subscript c) is

$$\rho_c = \frac{\frac{\sigma^2}{2} - G}{\sigma_r \sigma_{\bar{\eta}} \exp\left[\sigma_r^2/2\right]}. \tag{2.5}$$

We see that the threshold value increases with volatility, decreases with average resource growth rate, and decreases in magnitude with variance in either growth rate or resource variance. Progressive assignments are manifestly less feasible in more volatile and heterogeneous societies. Similarly, the critical volatility marking the crossover from positive to negative average growth, denoted σ_c, can be determined by rearranging Eqn. 2.5 for σ. This

relationship is plotted in Fig. 2.2, and gives a theoretical benchmark volatility for sustainable growth in a population for a set of distribution and growth parameters. Critical volatility decreases as a society becomes more progressive.

We can directly compare progressive and regressive assignments by computing the ratio of growth rates between a population with population effective growth rate Γ', and its progressive counterpart, Γ'_p with distribution parameters with superscript p. Set $\Gamma' = \Gamma'_p$ and rearrange for the ratio of mean growth rates between the population G^p/G. In the simplest case, where I assume identical initial population conditions up to the variance in growth rates with opposite correlation coefficients, $\rho^p = -\rho > 0$, the volatilities cancel out and the ratio is given by

$$\frac{G^p}{G} = 1 + \frac{1}{G}\rho\sigma_r(\sigma_{\bar{\eta}} + \sigma^p_{\bar{\eta}}) \exp\Big[\sigma_r^2/2\Big]. \tag{2.6}$$

Thus, a progressive arrangement must have a larger average growth rate in order to achieve a population effective growth rate equal to its regressive counterpart. This difference is greater in populations with stronger initial inequality, σ_r, and means that a higher progressive growth rate, G^p, is required to overcome larger growth rate variances, $\sigma_{\bar{\eta}}+\sigma^p_{\bar{\eta}}$. Fig. 2.2 demonstrates this equation's agreement with population Monte-Carlo simulations in which r_0 is normally distributed.

2.2.2 Fokker Planck Solution

So far, I have shown that the instantaneous population average growth rate collects a correction from covariances between resources and growth rates. This introduces a trade-off between managing inequality in the short run and maximizing overall growth which can only be resolved, at each time interval, if progressive assignments also result in higher aver-

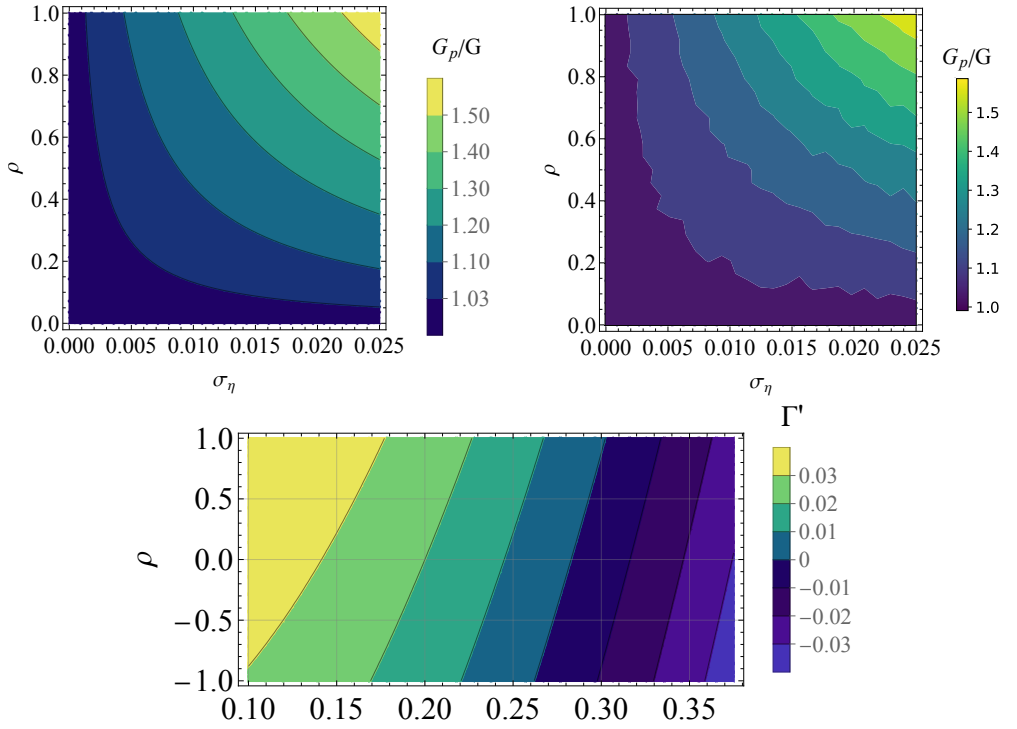

Figure 2.2: Parameter spaces for progressive/regressive growth rate ratios across variances (top) and growth rate regimes across correlation and volatility values (bottom) *Upper Left:* Contour plot of G^p/G for correlations and growth rate variances with condition $\sigma_{\bar{\eta}} = \sigma_{\bar{\eta}}^p$. *Upper Right:* A Monte-Carlo simulation comparing the correlated average resource growth rates of populations with identical parameters up to a change in sign of ρ agrees with theory. To stay competitive, the progressive society requires a larger growth rate as either parameter increases. *Lower:* The effective growth rate for $\mu_0 = .04, \sigma_r = 1, \sigma_{\bar{\eta}} = .025$ transitions from positive to negative along the correlation axis at high volatility ($\sigma \approx .275$). Γ'

age growth rates.

But what happens over the longer term, as growth rate fluctuations persist? To answer this question, consider the solution to the Fokker-Planck equation (FPE) for the time-dependent probability of resources for each agent. The FPE for the wealth distribution $P(r,t)$ for geometric Brownian motion is given by

$$\frac{\partial P}{\partial t} = -\frac{\partial}{\partial r}\left[\gamma r P(r,t)\right] + \frac{\partial^2}{\partial r^2}\left[\frac{\sigma^2}{2} r^2 P(r,t)\right]. \tag{2.7}$$

Dynamics of a Homogeneous Population

The solution with initial conditions $r(0) = r_0$ (a delta function at the individual level) is a lognormal distribution [230] with time-dependent mean and variance. This can be re-written as a Gaussian distribution of $\ln r$ as

$$P(\ln r, t|\gamma, \ln r_0) = \frac{d\ln r}{\sqrt{2\pi\sigma^2 t}} \exp\left[- \frac{(\ln r - \ln r_0 - \gamma t)^2}{2\sigma^2 t} \right], \tag{2.8}$$

where at early times, the expected resources are given by the linear equation $\ln \bar{r}(t) = \ln r_0 + \gamma t$, which asymptotically approach γt. The variance of log resources increases linearly in time, $\sigma^2 t$, which shows that individual wealth also becomes more uncertain, and population wealth becomes more unequal as a function of time. The critical volatility at which growth ceases is, as before, given by $\sigma_c = \sqrt{2\bar{\eta}}$. Consider Eqn. 2.8 to be the *homogeneous* population solution.

Changing variables and working explicitly with time-averaged growth rates yields further insight into the dynamics. Dividing each factor of the exponential by t, and performing the variable transformation $\frac{1}{t}\ln\frac{r}{r_0} \to \phi, t d\phi = d\ln r$ yields the distribution of growth rates

$$P(\phi, t|\gamma, \sigma) = \frac{d\phi}{\sqrt{2\pi\sigma^2/t}} \exp\left[- \frac{(\phi - \gamma)^2}{2\sigma^2/t} \right]. \tag{2.9}$$

This shows explicitly that the statistics of the time-averaged growth rates are simpler than those of resources. The average growth rate becomes stationary at long times as the mean remains constant, while growth rate fluctuations over the population decrease in time as t^{-1}, converging in distribution to a delta function on a time scale $t >> \sigma^2$.

This suggests a simple picture emerging at the population level; however, covarying the growth rates among agents in the population will introduce some complications. The most direct route to assess these effects follows by positing Gaussian distributions on growth rates, from the asymptotic behavior of growth rate distributions [23, 235, 249], and initial log resources across the population, from the solution to the FPE.

To obtain the dynamical solution for the distribution of resources in the population, I convolve Eq. 2.8 with a bivariate static Gaussian as described in 2.2.1. If the volatilities are held constant such that the marginal $P(\gamma)d\gamma = P(\bar{\eta} - \sigma^2/2)d\bar{\eta}$, we can perform the convolution over effective growth rates via the equation

$$P(\ln r, t) = \int P(\ln r, t|\gamma, \ln r_0)P(\ln r_0, \gamma)d\ln r_0 d\gamma. \tag{2.10}$$

The calculation is a straightforward Gaussian. The growth rate terms of the bivariate Gaussian distribution equation can be factored out of the resource integral by completing the square. The resource and growth rate integrals are then successively evaluated using Gaussian integration, leading to

$$P(\ln r, t) = \frac{d\ln r}{\sqrt{2\pi\Sigma^2}} \exp\left[- \frac{(\ln r - \ln \mu_0 - \Gamma t)^2}{2\Sigma^2}\right], \tag{2.11}$$

where $\Sigma^2 = \sigma_r^2 + (\sigma^2 + 2\sigma_r\sigma_{\bar{\eta}}\rho)t + \sigma_{\bar{\eta}}^2 t^2$. The covariance matrix is thus deduced as

$$K_\Sigma = \begin{pmatrix} \sigma_r^2 & 0 & \sigma_r\sigma_{\bar{\eta}}\rho t \\ 0 & \sigma^2 t & 0 \\ \sigma_{\bar{\eta}}\sigma_r\rho t & 0 & \sigma_{\bar{\eta}}^2 t^2 \end{pmatrix}. \tag{2.12}$$

Consider Eq. 2.11 the *heterogeneous* population solution. The quadratic expression in time for variance is characterized by two different timescales, $t_{c1} = \sigma_r^2/(\sigma^2 + 2\rho\sigma_r\sigma_{\bar{\eta}})$ and

$t_{c2} = (\sigma^2 + 2\rho\sigma_r\sigma_{\bar{\eta}})/\sigma_{\bar{\eta}}^2$. The natural timescale in this analysis is years [197], such that *early* times are on the order of a few years, *intermediate* over a few decades, and *long* timescales are several decades to centuries. Effects over long timescales can thus be thought of as generational, although I do not incorporate life cycle dynamics in this analysis. Annual growth rates are typically on the order of a few percent a year ($\gamma \simeq 10^{-2}$/yr), whereas log resources vary in the population on a typical scale of $\delta \ln r \simeq 10^1$, so the magnitude of $\sigma_{\bar{\eta}}$ is naturally an order of magnitude smaller than σ_r. This results in $t_{c1} < t_{c2}$, producing three distinct dynamical regimes for the population variance and hence inequality dynamics. For early times $t < t_{c1}$, the variance of resources in the population is given approximately by its initial condition $\Sigma^2 \simeq \sigma_r^2$. For intermediate times $t_{c1} < t < t_{c2}$, $\Sigma^2 \simeq (\sigma^2 + 2\rho\sigma_r\sigma_{\bar{\eta}})t$. The intermediate regime introduces the explicit decrease in inequality from a progressive assignment, so long as $\sigma^2 < 2\rho\sigma_r\sigma_{\bar{\eta}}$. However, this benefit is short-lived, as for later times $t > t_{c2}$, $\Sigma^2 \simeq \sigma_{\bar{\eta}}^2 t^2$, and variance invariably explodes due to fluctuations in growth rates. Note that the long-time regime does not require biases between resources and growth rates; it persists under the weaker conditions of a finite variance in growth rates and appears at earlier times for larger growth rate variances.

Finally, we observe that the asymptotic population averaged resource growth rate distribution simplifies to a Gaussian

$$\lim_{t\to\infty} P(\phi, t) = \frac{d\phi}{\sqrt{2\pi\sigma_{\bar{\eta}}^2}} \exp\left[-\frac{(\phi - \Gamma)^2}{2\sigma_{\bar{\eta}}^2} \right], \tag{2.13}$$

in contrast to the case of a homogeneous population, reflecting the Gaussian distribution of growth rates in the population. Over the long term, population dynamics are thus dominated by fluctuations in growth rates, leading to wider inequality the larger the growth rate variance.

This effect introduces a general mechanism that can account for increases in inequality

that occur over intermediate and long timescales [95], and incorporates findings that empirical growth rates are heterogeneous and correlate with wealth over time [77, 110, 125]. It also asks that we shift focus from the dynamics of wealth itself to the variations of growth rates as the drivers of long-term inequality.

2.3 Measuring the Dynamics of Inequality

This section analyzes the detailed dynamics of inequality using two standard metrics. First, the Gini coefficient, G_{ini}, is the most common metric measuring inequality within a population. It ranges from $G_{\text{ini}} = 0$ for equally distributed resources, and trends towards 1 as a society becomes maximally unequal. For lognormally distributed resources in the infinite population limit, $G_{ini} = \text{Erf}\big[\Sigma/2\big]$. Second, I examine the social cost of high growth on inequality, a topic of interest and debate in the social sciences [56, 164]. Using the coefficient of variation (CV), denoted c_v, I compare how quickly the standard deviation of a distribution of form 2.11 increases relative to its mean [19]. The dynamics are given by $c_v = \Sigma/(\ln \mu_0 + \Gamma t)$, and should be governed over long times by the linear time ordering of both the mean and standard deviation terms. A more equitable economy would decrease c_v; growing without increasing inequality. In the egalitarian limit, $c_v = 0$ as every agent has the same resources, and variance is zero. Conversely, an exploding c_v in the positive or negative direction indicates increasing inequality, marked by positive or negative aggregate growth respectively. From Eq. 2.8, without dynamical redistribution measures, societies will invariably become maximally unequal as multiplicative dynamics are dominated by fluctuations in earnings [213].

I develop the dynamics of homogeneous and heterogeneous populations by comparing time-series trajectories of c_v and G_{ini} across several values of correlation coefficient and volatility. Fig. 2.3 demonstrates the impact of heterogeneity on inequality (Gini), and its relative effects on growth (CV). In homogeneous populations, volatility spurs inequality from

uniform initial conditions causing c_v to initially increase, peak, then decrease as resources increase in $t^{1/2}$ faster than the population standard deviation. The quantity G_{ini} asymptotically approaches 1 in all parameter cases. Heterogeneity causes c_v to increase across all cases with asymptotically constant behavior as the time leading terms of the numerator and denominator of c_v cancel. Heterogeneity causes G_{ini} to generally increase more rapidly. The differences in progressive and regressive G_{ini} are negligible after intermediate times, suggesting that the initial configuration has little effect on its long-term macroscopic dynamics. Volatility strongly determines the value of c_v for homogeneous and heterogeneous populations, while only significantly affecting the behavior of G_{ini} in the homogeneous case.

The dynamics of G_{ini} are thus straightforward, but the dependence of c_v on both growth rates and dynamical variance produces interesting dynamics. To probe the dynamics of c_v, I produce quasi-phase diagrams displaying $\partial_t c_v$ for a continuous spectrum of volatilities over time. Fig. 2.4 demonstrates these dynamical regimes for the homogeneous and heterogeneous progressive and regressive cases. It shows the magnitude of the initial increase in c_v scales with σ, as does the magnitude of the eventual decrease. The eventual decrease is weaker in regressive populations, occurring only at medium volatility values, and is nonexistent for any volatility value of the progressive population. Furthermore, they reveal a divergent regime past the critical volatility, where resources decrease while inequality continues to increase.

2.4 Discussion

These results show that growth rate fluctuations pose a general challenge to any proposed mechanism to address widening inequality that must be addressed independently of volatility in single-agent trajectories. Specifically, dynamically balancing growth rates is required to control long-term inequality dynamics. Methods seeking to both reduce growth rate fluctuations and preserve growth require a theory for the origins of wealth growth rate

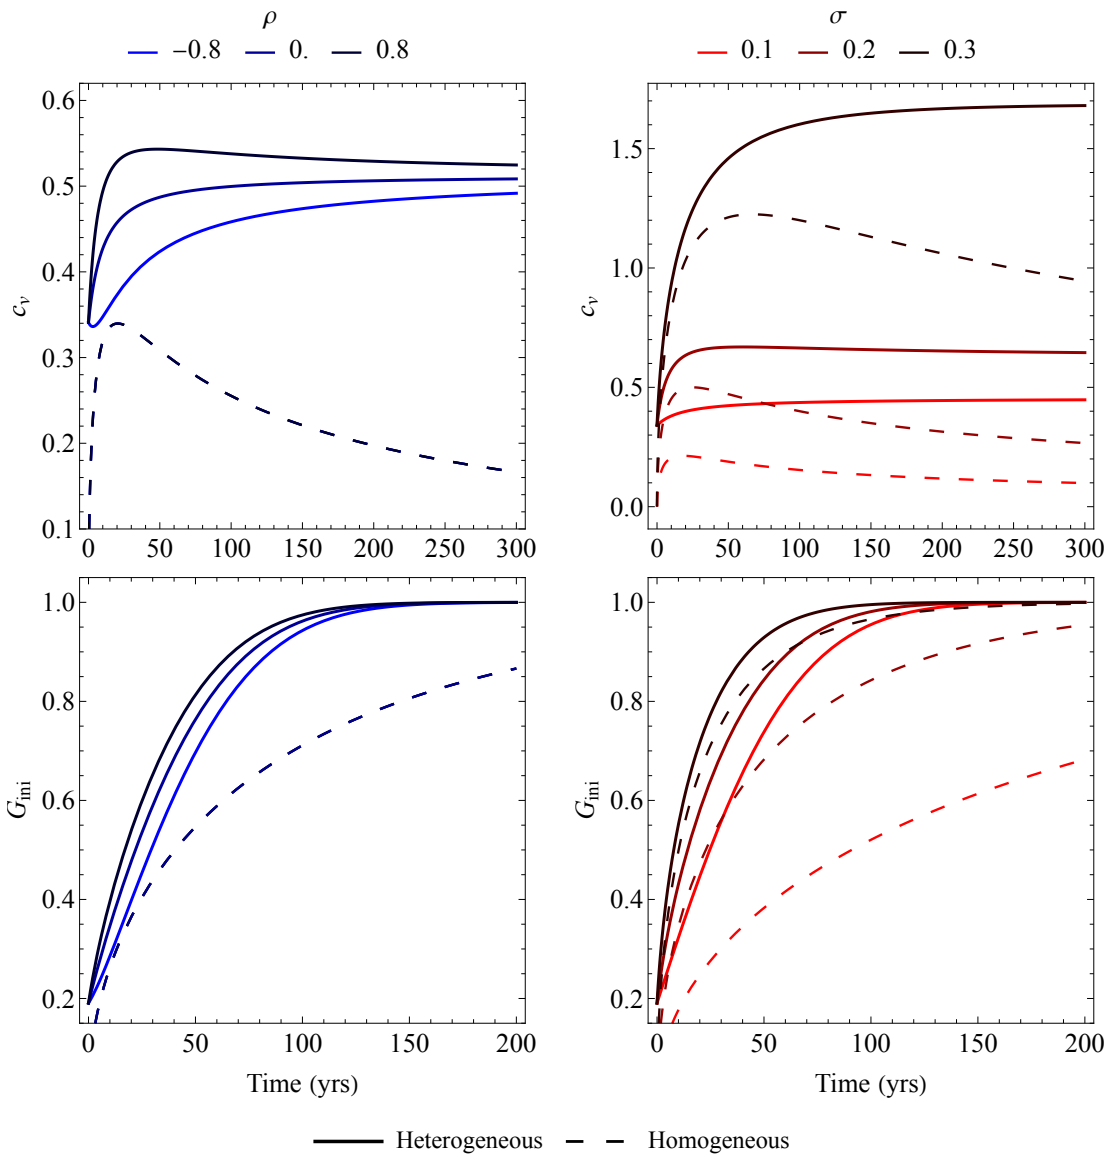

Figure 2.3: Trajectories of inequality metrics in various parameter regimes; $\sigma = .15, \sigma_r = .341, \ln \mu_0 = 1, G = .06, \sigma_{\bar{\eta}} = .025$. See legend for line type mapping. *Left:* Plots for low volatility populations with different correlation coefficient values. For $\rho > 0$, c_v jumps higher than in an uncorrelated society, while c_v temporarily decreases than increases for $\rho < 0$. All heterogeneous trajectories approach similar constant values, while the homogeneous population initially increases, then decreases over long times. At early times, G_{ini} increases more quickly for $\rho > 0$. In all cases, variances in growth rates drive G_{ini} to 1 more rapidly than in homogeneous populations. *Right:* Regressive populations ($\rho = .75$) at different values of volatility. A higher volatility causes a more rapid increase in G_{ini} in both population configurations, and volatility positively influences the peak and asymptotic values of c_v. Progressive assignments hurt long-term growth while marginally affecting long-term inequality dynamics, while heterogeneities play a dominant role in accelerating the emergence of inequality.

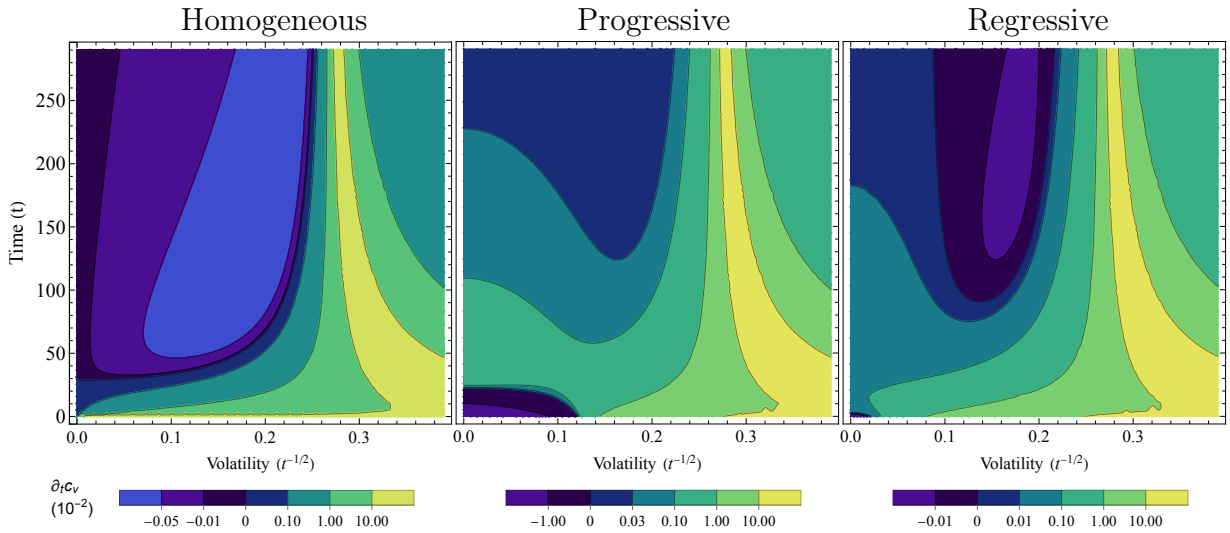

Figure 2.4: Contour plots for $\partial_t c_v$ in identical populations to Fig. 2.3, with $\bar{\eta} = .035$. At low volatilities, homogeneous c_v initially increases, then decreases with magnitude increasing in σ. For $\rho > 0$, this behavior persists following initial times of high increase in c_v. The initial decrease observed for $\rho < 0$ only occurs at low volatility, and c_v does not decrease in later times, as demonstrated in Fig. 2.3. Past the homogeneous critical volatility, given by $\sigma_c \approx .265$, c_v blows up as the negative growth rate drives average log resources through zero to a negative value, and is seen by the bright yellow bands. In this regime of instability, resources continue to decrease as inequality grows.

statistics. Extending the stochastic growth models commonly used to include environmental opportunities and agent-based learning, in ways analogous to gambling [136] and portfolio theory [123], naturally produces situations where individual growth rates become distinct and history-dependent but also where average growth rates can be maximized and variances minimized over time [23, 60, 136]. In this sense, growth rate redistribution takes the form of investment in learning across the population, through for example education, which over time reduces disparities in growth rates. In such a dynamical picture, initial correlations will also take on dynamic properties in ways stated here. In recognizing recent research on the effects of heterogeneous and dynamical growth on distributions of wealth [1, 42, 124, 179, 230], we seek a theoretical framework for wealth dynamics that both complements these phenomenological approaches and incorporates strategic agent behavior and adaptation to changing circumstances in statistical environments. I will derive and explore this model in Chapter 3, and will discuss its implications for heterogeneity-induced inequality explored in this chapter.

In summary, I explored the effects of fluctuations in wealth growth rates and wealth statistics in the temporal development of inequality in a population of heterogeneous agents. I set up the general problem in the context of multiplicative random growth and derived closed-form expressions for how the population inequality changes over time, thus identifying the most important parameters at early and later times. I found that population variances in growth rates, their correlations with resources, and individual temporal volatility are the primary drivers of inequality in general situations. The sustained presence of variance in growth rates in a population produces a dominant effect on long-term inequality. While redistribution methods, such as those explored by mean-field or other more complex network exchange models [21, 39, 97] reduce the dynamical impacts of volatility, I have shown that directly addressing growth rate variance in a population is a more fundamental and general requirement for arresting or reversing increases in wealth inequality.

CHAPTER 3

ADAPTIVE GROWTH IN STOCHASTIC ENVIRONMENTS

As discussed in Chapter 2, there remain general questions about how societies can promote long-term growth while controlling or mitigating inequality. Compared to other configurations, assigning high growth rates to agents with a lot of resources increases average growth while driving up inequality on intermediate timescales. In general, heterogeneity in growth rates increases inequality on a higher order in time than volatility alone. Wealth redistribution addresses the state of resource inequality by reallocating resources [197]. However, an analogous mechanism is needed to reduce inequality in growth rates, and therefore inequality in the average dynamics. To derive such a mechanism, we first need a theory for the origin of growth rates in statistical environments.

Growth and inequality can be more holistically understood in terms of the agent's behavior relative to an environment. This chapter reframes the study of these emergent properties in terms of an adaptive process derived from theory for optimal agent behavior[1].

3.1 Introduction

In much of the growth literature, agents representing individuals or households (often with life cycles), grow or lose wealth through a multiplicative (geometric) stochastic process. This modeling choice is well supported empirically and introduces a number of key parameters as an agent's resources (or wealth), r, evolve exponentially with mean growth rate (over time), γ, fluctuate with standard deviation (volatility), σ [23, 39, 188] and vary across individuals of a population with standard deviation σ_γ [95, 140]. These parameters describe the statistical dynamics of wealth in heterogeneous populations and the emergence of inequality across various timescales.

1. This chapter is adapted from Kemp, J. T., & Bettencourt, L. M. (2023). Learning increases growth and reduces inequality in shared noisy environments. *PNAS nexus, 2(4), pgad093.*

Recent developments in cognitive and ecological sciences can provide some valuable insights into the stochastic dynamics of agent behavior [240]. Researchers using noisy decision-making models to explore child and adolescent development have recently rethought the process of human learning in terms of acquiring information through (active and passive) interactions with a knowable, but stochastic external environment [52, 122, 247]. Similarly, ecologists have formulated natural selection, the process through which a genotype optimally leverages its environment's structure to maximize population growth (fitness), as a (Bayesian) optimization process [27, 43, 84, 89, 152]. As mentioned in Chapter 1, biophysicists have come to similar conclusions [152, 204] These approaches describe (individual or collective) agent optimal choices as the result of information they obtain in a noisy, but knowable environment, with information dynamics that are fundamentally Bayesian. This connection between optimal intertemporal decisions, information, and fitness (growth) was previously explored as a mathematical formalism to optimize betting and portfolio investment returns [60, 137]. However, its applications to human behavior and population dynamics suggest it serves a suitable basis for the general statistical mechanics of wealth growth and inequality [27].

In this chapter I unify these approaches to develop a statistical dynamics of growth and inequality in a population of strategic, adaptive agents, where the growth rates result from investing and learning in a stochastic environment.

In this approach, heterogeneous agents invest in sequential, stochastic environmental events based on signals they perceive as they go, and grow their wealth based on the quality of their predicted allocations. By exploring this mechanism of (optimal) information-driven growth in the context of population dynamics we obtain a better understanding of how wealth growth and disparities originate from differences in agent knowledge and adaptive behavior. More broadly, this work adds a new dimension to the study of wealth dynamics [2]

2. and as I will explore later, interactive dynamics

that more fundamentally links inequalities between wealth, growth, and agents' subjective characteristics, such as their present knowledge, their singular life-course experience, and the quality of their knowable environment, e.g. in terms its opportunities expressed as statistical rates of return on investments.

Our approach treats both resources and information as dynamically coupled quantities. To model information dynamics, I show that learning in the joint space of environmental states and agents' signals is developed optimally in terms of Bayesian inference, translating a maximization of the predictability of environmental states into that of resource allocations and growth.

I finish this chapter by exploring the general consequences of learning a shared environment on the statistics of information and resources, and discuss the consequences for the role of general education and training on population dynamics and its potential to reverse long-term wealth inequality [140].

3.2 Theory and Modelling of Information-Based Growth

I start by deriving a general theory of growth rates in terms of informational quantities. Here, information means an agent's predictive knowledge of event probabilities in a noisy environment. Agents seek to maximize the growth of their resources over time by investing in a set of possible events in their environment using their individual knowledge. Agent's knowledge is subjective, as it is formed by the agent's own experience, model of the world and expectations ("beliefs"), which are assumed here not to be shared or compared with other agents. The agent's beliefs are adjusted by observing environmental outcomes in time through an iterative process of (Bayesian) learning. After developing the general framework, I illustrate these dynamics using a multinomial model of discrete environmental states and choice, for which I derive closed-form expressions for the average resource growth rate and volatility in terms of information-theoretic quantities. I will then identify the general cir-

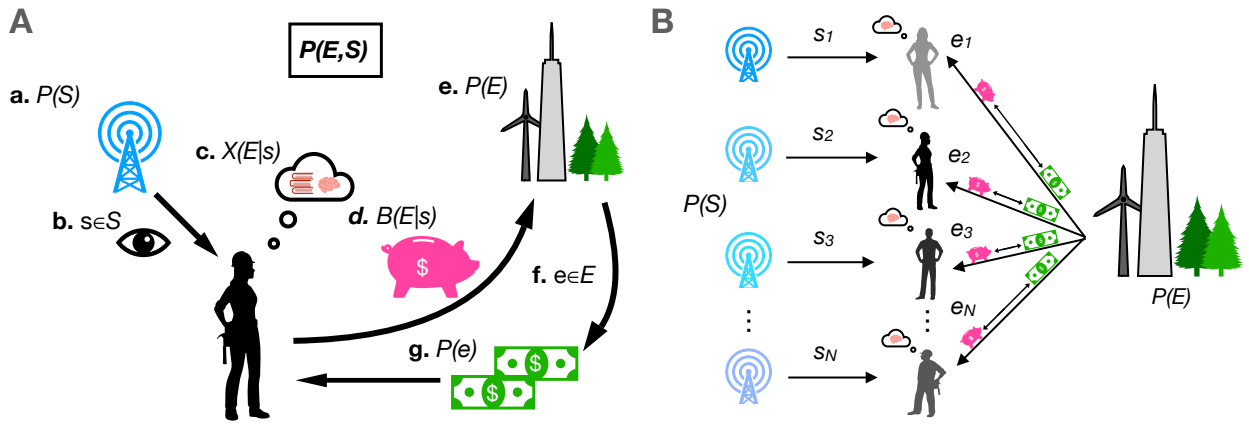

Figure 3.1: General dynamics of adaptive growth: Agents obtain resources from their environment based on the quality of their information. **A.** At each time step, **a**. the agent's private channel (memory, senses) outputs a signal $s \in S$ with probability $P(s)$. **b**. The agent observes the state s, **c**. and consults their belief for the conditional outcome probability of the environment, $X(E|s)$. **d.** The agent makes proportional resource allocations on all possible outcomes $B(E|s)$. **f,e.** The true event $e \in E$ is revealed from the environment with probability $P(e)$, and **g.** the agent receives a payout proportional to the marginal probability of e. **B.** In a population simulation, N agents independently sample private signals and invest in events sampled from the same environment.

cumstance when this learning process dynamically attenuates inequality in resource growth rates across populations.

3.2.1 Growth from Information

Consider a population of $i = 1, ..., N$ agents, each with initial resources r_i that can be (re)invested into the set of outcomes of their environment to generate returns. The agents have access to a private signal $s \in S$, which they use as a predictor to invest resources in events $e \in E$ generated by their environment. The set of signals and events are described by the joint probability distribution, $P(E, S)$ with marginals $P(E)$ and $P(S)$.

At every time step, each agent observes its own signal s, and allocates resources r on events following a vector $B(E|s)$, such that $\sum_e B(e|s) = 1,\ e \in E$. As the event e is revealed, the agent is awarded returns, w_e for the fraction of resources invested in the correct outcome, $B(e|s)r_i$. After n steps, the agent's total resources (wealth) are

$$r_n = r_i \prod_{j=1}^{n} B(e_j|s_j) w_{e_j} = r_i \prod_{s,e} \left[B(e|s) w_e\right]^{W_{s,e}}, \tag{3.1}$$

where $W_{s,e}$ is the number of occurrences ("wins") of s, e. By the law of large numbers, $\frac{W_{s,e}}{n} \to P(s,e)$ as $n \to \infty$. It follows that the average growth rate of resources over large n steps is

$$\gamma_i \equiv \frac{1}{n} \log \frac{r_n}{r_i} \approx \sum_{e,s} P(s,e) \log[B(e|s) w_e]. \tag{3.2}$$

Kelly showed that the maximal growth rate as $n \to \infty$, obtained by maximizing Eq. (3.2) with relation to $B(E|S)$, results in an allocation mirroring the conditional probability, $B(E|S) = P(E|S)$. This maximum growth rate is the mutual information, $\gamma_{max} = I(E,S)$ when the odds are "fair", $w_e = 1/P(e)$ [137].

Typical agents do not start out with perfect knowledge. In this case, agents must invest resources using their best estimate for the conditional probability, $X(E|S) \neq P(E|S)$. Then, their resource growth rate will be lower than the maximum. This can still be written in terms of informational quantities as the Kelly growth rate (SM 1),

$$\gamma = I(E;S) - \mathrm{E}_s\big(D_{KL}\big[P(E|s)||X(E|s)\big]\big). \tag{3.3}$$

where E_s is an expectation value over the states of the signal, and $D_{KL}\big[P(E|s)||X(E|s)\big] = \sum_e P(e|s) \log \frac{P(e|s)}{X(e|s)} \geq 0$ is the Kullback-Leibler divergence, expressing how similar the two distributions in its inputs are. This general result shows that agents with better information will experience greater resource growth rates, as long as they invest optimally [5]. These compounding dynamics are illustrated in Fig. 3.1. we also see that this setup allows us to consider agents with different knowledge, corresponding to skill heterogeneity within a population. I will discuss other general issues of innovation and structural position as I introduce learning in populations below.

Note that in reality, people may be in debt, typically leading to a negative component of their growth rate due to interest payments. I do not consider this situation here, except to point out that if such a component is constant it does not affect my analysis. If, however, the loan rate can be reduced via better information, it will add another dimension to the optimization of the overall growth rate.

I will now illustrate these general results using a specific, stationary multinomial model. While the theory is developed for general environmental dynamics, its limitation to a stationary environment will allow us to derive quantities of mean growth rate and volatility, familiar to GBM, in closed-form and establish the parallels to most wealth-growth models. This model will allow us to then illustrate and simulate the population dynamics of growth and inequality among agents with heterogeneous information. Later, I will also show how agents can improve their information optimally over time through a process of iterative Bayesian learning.

3.2.2 Multinomial Choice Model

Consider the space of signals S and environmental states E of equal size l, with outcomes $s, e \in 1, ..., l$ and degenerate, multinomial conditional probability

$$P(e|s) = f(p, l) = \begin{cases} p & \text{if } s = e \\ \frac{1-p}{l-1} & \text{if } s \neq e, \end{cases} \tag{3.4}$$

where $0 < p < 1$ is the binomial probability of guessing the correct environmental outcome. For simplicity, I assumed that the probability of a correct guess is independent of l. The distribution has uniform marginals, $P(e) = 1/l$ and $P(s) = 1/l$, for all signals and events, such that $P(s|e) = P(e|s)$ via Bayes' rule.

With these choices, we can derive expressions for the relevant informational quantities in closed-form. The mutual information between an agent's signals and environmental outcomes

is then $I(E;S) = \log l + p \log p + (1-p) \log \frac{1-p}{l-1}$ (APP A.14). As the simplest illustration, for a binary choice, $l = 2$, the first term gives 1 bit of entropy of the environment and the remaining terms give the conditional entropy, expressing how well an agent could know the environment given their signal. In the limit $p \to 1$, the signal gives agents perfect knowledge of $P(E)$.

So far I considered that the agent has perfect knowledge of the joint distribution of the signals and the environment. When this is not the case, we can write a parametric expression of the agent's ignorance in terms of an estimated binomial probability $x \neq p$. The agent's likelihood model of the conditional probability is then $X(e|s) = f(x, l)$. The divergence term of Eq. 3.3 becomes the divergence between $f(p, l)$ and $f(x, l)$ averaged over all signals, $\mathrm{E}_s\big[D_{KL}\big] = p \log \frac{p}{x} + (1-p) \log \frac{1-p}{1-x}$. Subtracting the mutual information by this term yields the agent's actual growth rate under imperfect information as (SI A.2.1)

$$\gamma = \log l + p \log x + (1-p) \log \frac{1-x}{l-1}. \tag{3.5}$$

This expression is plotted in Fig. 3.2A as a function of x for various l values and fixed p. We see that increasing the size of the event space, l, reduces the probability of any individual outcome, increasing the payouts and the Kelly growth rate. The maximal growth rate is obtained when $E_e[D_{KL}] \to 0$, when $x \to p$. Conversely, $\gamma \to 0$ when $p \to 1/l$, indicating the signal and the environment have become statistically independent.

Treating γ as the expected resource growth rate, the volatility is calculated as the second moment of the growth process. The volatility squared (variance) is given as (APP A.2.1)

$$\sigma^2 = p(1-p) \log^2 \frac{x(l-1)}{1-x}. \tag{3.6}$$

This expression is shown in Fig. 3.2C. The volatility vanishes in the limit $x \to 1/l$, corre-

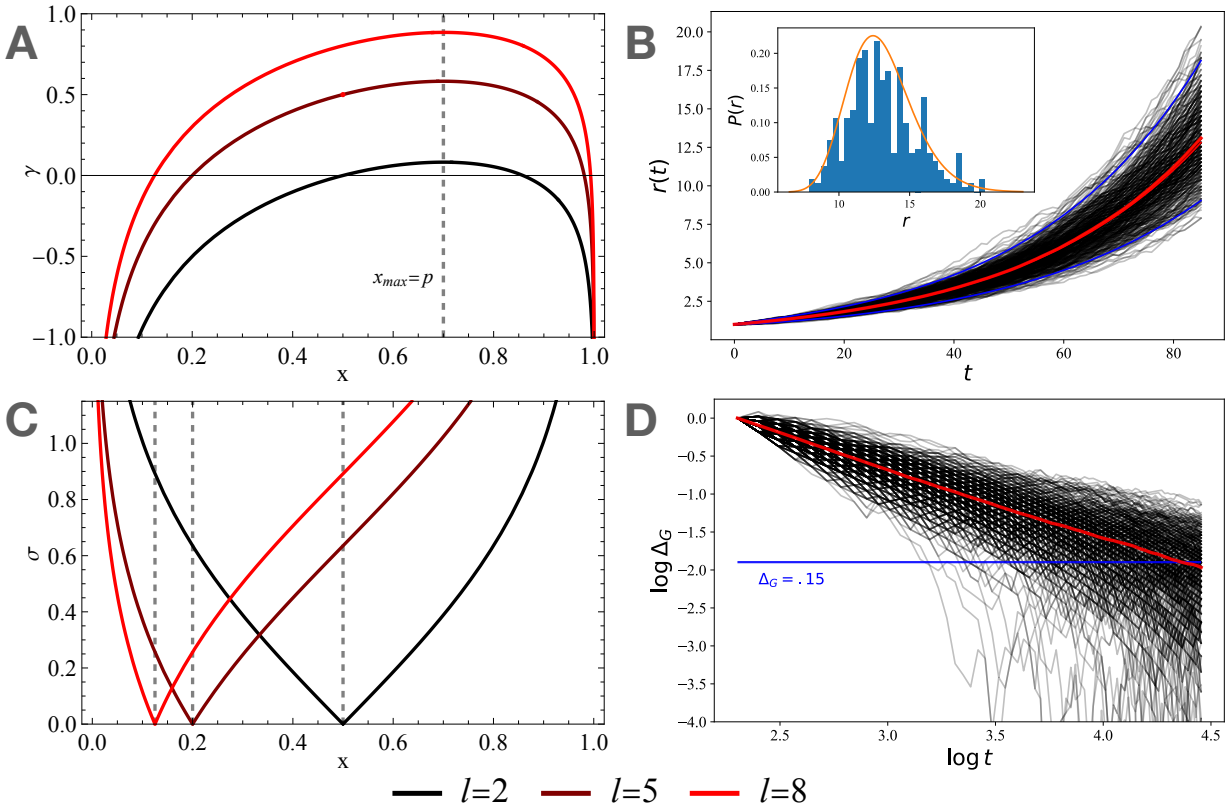

Figure 3.2: Growth dynamics curves and MC simulations of growth process without learning. The growth rate and volatility are computed analytically for a discrete multinomial environment, reproducing the limit of GBM dynamics. **A**. For $p = .7$, the growth rate maximizes at $x = .7$, decreases as x diverges from p, and scales with l. The parameter $l = 2$ provides a realistic range of average growth rates. **B**. Monte Carlo simulations with $N = 388$ homogeneous agents, all with $\gamma(x) = .03$ and $r_0 = 1$. The expected mean (red) predicted by $\gamma = .03$ and population mean (orange) overlap in value. The blue lines represent the 95% confidence interval. *Inset*: The resource histogram is fit to a log-normal distribution of the same growth and volatility parameters. **C**. Volatility is minimized at $x = l/2$ and increases monotonically in either direction. Volatility increases more rapidly at higher values of l. **D**. Over time, $\Delta_\gamma \to 0$ as agents' growth rates approach the Kelly growth rate. The average agent converges to within 15% the expected mean at $t \approx 80$.

sponding to when agents invest indiscriminately with equal probability in all possible event types. A larger l increases the magnitude of the growth rate, but also the volatility. The volatility is highest when $p \to 1/2$ and the environment is most uncertain. In any case, the agents feel surest of the outcomes when $x \to 0$, or $x \to 1$.

Kelly's formulation describes the average growth rate of resources over a large number of discrete investments [137]. To derive a growth process in time, I average over ω bets per unit time, such that $\Delta t = 1/\omega$ is the interval of time between investment periods. Returns at time $t + \Delta t$ are then the mean of all investment returns earned in the time interval $[t, t + \Delta t]$. In the limit $\omega \to \infty$, as the agent makes continuous allocations, $r_n \to r(t)$ and γ describe the average growth rate. Consider $t \approx 10^{-2} yr$ (i.e. 1% a year) so that the simulated results are comparable to previous work based on yearly growth rates of the order of a few percent. Volatility is reduced $\sigma_t = \sigma_n/\sqrt{\omega}$ as fluctuations are averaged out in each time step (SM 10).

Fig. 3.2C demonstrates the two investment regimes for each value of γ, where the growth rate maps to either high or low volatility depending on the value of x. Investments with $x > p$, which are described as *aggressive*, overestimate the dependence between the signal and environment. Under this condition, agents invest relatively more on diagonal outcomes and experience large gains or losses resulting in higher volatility. With $x < p$, which is denoted *conservative*, agents underestimate p and distribute their wealth more equally across all outcomes, resulting in less volatility. Agents can also experience $\gamma = 0$ at two values of x: In the trivial limit, as $x \to 1/l$, signals and agent investments become statistically independent. The other trivial case can be solved for numerically when $\gamma = 0$.

With given x independent of time, the dynamics reduce to the well-known behavior of geometric Brownian motion (GBM) with drift. Fig 3.2B shows the dynamics of a population of agents with homogeneous (non-time dependent) parameters evolved using a Monte-Carlo simulation. In this particular situation, mean population resources grow with $\langle r(t) \rangle =$

$\frac{1}{N}\sum_i r_i(t) = \exp[\gamma t]$, in agreement with [23].

I also demonstrate that the time-averaged growth rate of resources converges to the Kelly growth rate over many allocations. Fig 3.2D shows the asymptotic convergence of the normalized difference of averaged growth rate for individual agents $\Delta_G = (\gamma - G)/\gamma \to 0$, where $G = \frac{1}{t}\ln\frac{r(t)}{r(0)}$ (black) and population-averaged growth rate, $\langle G\rangle = \frac{1}{N}\sum_i G_i$ (red).

I have thus far considered x as a static variable and explored the dynamics of resources when $x \neq p$ in a stationary environment. To converge to maximal growth rates, however, it is necessary that agents can estimate the correct event properties, given their signals, a situation to which I now turn.

3.2.3 Dynamical Growth Rates from Bayesian Inference

Realistic agent trajectories are dynamical, reflecting investment allocations that are history-dependent and result from the cumulative knowledge of each agent's past experience [15, 23]. Agents must then improve their information about the environment by updating their model of the conditional relationship of $S|E$ with each observation. In the absence of other random processes, this learning task is optimally achieved in terms of sequential Bayesian inference [17, 61]:

$$X_n(e|s) = AP(s_n|e_n)X(e_n) = \left[\Pi_{i=1}^{n}\frac{P(s_i|e_i)}{P(s_i)}\right]X(e), \tag{3.7}$$

where the normalization $A = \left(\int de_n P(s_n|e_n)X(e_n)\right)^{-1}$. I also take the prior probability, $X(e_1) = X(e)$, because I am assuming that the environment is stationary or at least slowly changing relative to agents' learning rates.

Then, Bayesian inference converges $X(E|S) \to P(E|S)$, decreasing the information divergence over long times. Through interactions with the environment, the agent optimally gathers information [27] as well as resources as demonstrated in Fig. 3.3A. Specifically, by minimizing the information divergence, learning agents maximize their resource growth over the long term.

This formalism allows us to start considering general aspects of innovation in heterogeneous populations, including issues of competitive advantage and structural positions in terms of agents' initial knowledge, models of the environment, and embedding within socioeconomic networks. Regardless of any of these elaborations, any learning model aspiring to optimal prediction must be Bayesian, as it is the single best way to incorporate observed data towards making predictions of future states of the environment [175] and maximizing long-term growth.

There is growing interest in incorporating learning agents in economics and other social sciences towards formulating models of more realistic "rational expectations" in intertemporal optimization problems [76]. At present, however, most of these approaches adopt simplified learning models, for example, based on least-squares minimization [76], which at best apply in particular cases, such as for Gaussian likelihoods. Consequently, we see a wide range of interesting opportunities in the social sciences for the adoption of more explicitly Bayesian frameworks, as has become increasingly common in psychology [105].

In the following subsection, I describe a parametric Bayesian inference scheme applied to the multinomial model via a Dirichlet prescription of conjugate priors [33], before I return to the general case to discuss issues of inequality in the light of learning.

3.2.4 Bayesian Dynamical Growth in the Multinomial Model

To illustrate these learning dynamics, we now return to the multinomial model of choice. I define the agent's likelihood function of a sample of the signal, $s|e$, as a categorical distribution with parameter vector $\beta = \{\boldsymbol{\beta}^1, ..., \boldsymbol{\beta}^l\} \in \mathbb{R}^l$, with each vector corresponding to an event and each component, β_s^e corresponding to a signal, event pair. The probability mass function is given by $P(s|e) = \prod_s (\beta_s^e)^s$, with normalization $\sum_s \beta_s^e = 1$. The conjugate prior distribution of E is given by a Dirichlet with hyperprior vector $\boldsymbol{\alpha} \in \mathbb{R}^l$, and distribution $P(e) = \alpha_e / A$, where magnitude $A = \sum_e^l \alpha_e / l$. This scheme is illustrated in Fig. 3.3B.

Set $\alpha_e = 1$ for all e so that the prior is uniform, for simplicity. I ensure the off-diagonal degenerate condition by setting $\beta_s^e = p_e$ for $s = e$, and for off-diagonal events, $e \neq s$, $\beta_s^e = \frac{1-p_e}{l-1}$, satisfying Eqn A.12. The binomial parameter describing the environment is then given by the average along the diagonal,

$$p = \frac{1}{l}\sum_{s}^{l} \beta_s^e, \qquad s = e. \tag{3.8}$$

An agent with imperfect information will have estimates for the parameters, $\tilde{\boldsymbol{\alpha}} \neq \boldsymbol{\alpha}$ and $\tilde{\beta} \neq \beta$, and posterior, $X(E|S, \tilde{\beta}, \tilde{\boldsymbol{\alpha}}) \neq P(E|S, \beta, \boldsymbol{\alpha})$. With each observation, the agent must update $X(E|S)$ via (APP A.2.2)

$$X(e|s) \propto \frac{m_{(-s)}^{(-e)}/\omega k + \tilde{\beta}_s^e}{M^{(-s)}/\omega k + 1}(n_{(-e)} + \tilde{\alpha}_e), \tag{3.9}$$

where $m_{(-s)}^{(-e)}$ and $n_{(-e)}$ are the total numbers of samples e, s and e excluding the current, and $M^{(-s)} = \sum_e m_{(-s)}^{(-e)}$ is the number of samples of s excluding the current. I also introduce an inference time, k, as a free parameter that weighs the evidence versus the prior, with units *time/update* such that t/k is unit-less. In the limit $k \to \infty$, the agent does not update their prior with new evidence. In the opposite limit, $k \to 0$, the agent ignores the prior and considers only the most recent evidence, and this becomes a maximum likelihood model.

During the inference process, the agent will break the degeneracy of their posterior as they infer each β_s^e individually. This is inconsequential though, as $x(t)$ can still be computed similarly to Eq. 3.8 at any time. The degeneracy of $P(E|S)$ permits us to reduce the dynamics of $X(E|S)$ to that of the diagonal probability $x(t)$, such that (APP A.2.2).

$$x(t) = \frac{pt/kl + x_0}{1 + t/kl}, \tag{3.10}$$

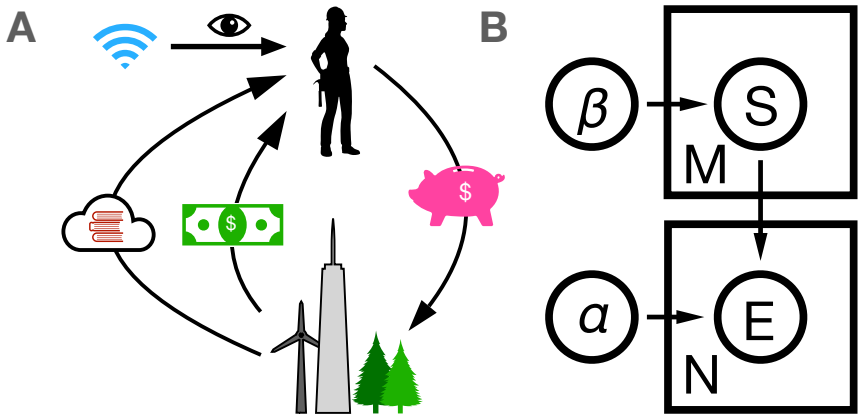

Figure 3.3: Schematic illustration of the learning process and the LDA algorithm. **A**. In addition to earning resources, the agent obtains information with each investment in the environment. **B**. Notation for the latent Dirichlet inference process. The agent is assigned prior parameters $\tilde{\alpha}, \tilde{\beta}$, corresponding to their belief for the distributions of E and S, which are updated based on event counts M, n respectively.

where x_0 is the agent's initial binomial probability parameter. This equation illustrates the core results of this approach, as the dynamics of the information of the agent's posterior determine the average dynamics of the growth rate via the functional, $\gamma[x(t)]$. Over many observations, the agent optimizes their guess, driving $X \to P$, minimizing their information divergence as $D_{KL}(P||X) \to 0$. The agent thus maximizes the average growth rate for their signal over time with a power law -1 in terms of the dimensionless inference parameter $\lambda \equiv t/kl$, at larger times $\lambda >> 1$. As previously mentioned, though, agents who have maximized information are still subject to the volatility of sample fluctuations. For the remainder of this chapter, I will study the effects of this learning process on the population dynamics of growth rates and wealth.

3.3 Population Effects of Information Dynamics

Having defined the dynamics of information and resources for single agents, we now explore the general dynamics of growth rate statistics in a heterogeneous population and its implications for long-term inequality. Mean growth rates can vary because of a number of

different factors. Particularly, agents have different initial conditions of knowledge, they experience different environmental stochastic histories, and they may have different models of the world in terms of their likelihood functions. I will now explore these sources of information heterogeneity and show that with a shared statistical signal, a population can reverse the (dominant) effects of heterogeneity on growth and inequality [140].

Consider the population variance of growth rates generally in terms of information-theoretic quantities, where $I_i \equiv I(E;S_i)$ and $D_i \equiv \mathrm{E}_{s_i}\big(D_{KL}\big[P(E|s_i)||X(E|s_i)\big]\big)$. The population variance is given as (APP A.2.3)

$$\mathrm{Var}_N[\gamma_i] = \mathrm{Var}_N[I_i] + \mathrm{Var}_N[D_i] - 2\mathrm{Covar}_N\big[I_i D_i\big]. \tag{3.11}$$

The first term is independent of any agent's imperfect knowledge or learning process and depends only on their model (likelihood) of the environment, given the agents' signals.

The second term expresses variance in the prior and different learning trajectories across agents. This term vanishes as agents learn their environment fully. It follows that these two sources of variance vanish only if every agent has the same model in a shared environment with the same statistics, and after every agent has had time to learn their environment.

These two terms also express formal distinctions between the familiar Keynesian formulation of intrinsic uncertainty versus risk in socioeconomic behavior. Agents cannot know *a priori* what type of uncertainty they are facing and must learn as best as they can from their experience. A misspecification of the agents' model of the world, via an incorrect likelihood function, will result in irreducible uncertainty and a lower growth rate than possible. In terms of communications theory, this situation effectively uses the environmental experience suboptimally, by picking a signal that does not maximize the channel capacity, as the largest possible mutual information between the agents' signal and events in the world [59]. On the other hand, risk in the sense of probabilistic events with a known distribution can be reduced (and better assessed) via the Bayesian inference process which builds the correct

risk model within a family of functions, by learning its parameters.

The third term is less familiar and arises in populations where the magnitude of the agents' information co-varies with agents' divergence from the environment. This may happen in reality when different (likelihood) models of the world co-exist in a population of agents, and when, in addition, less experienced agents with shorter learning histories, adopt preferentially some of these models. For example, a younger generation may have a better model of the world but less experience, creating a negative co-variance. Or a positive covariance may be generated if learners with a better model are encouraged to learn faster, and others discouraged, creating a kind of cumulative advantage in terms of better information and faster learning. Such situations may provide principled modeling strategies to better understand the success of *a posteriori* exceptionally successful individuals, and identify situations of competitive advantage in access to information and learning.

3.3.1 Population Effects in the Multinomial Model

This section illustrates how the inference dynamics happen in the context of the Multinomial model. I focus on agents with identically distributed signals i.e. with the same likelihood function and a shared environment, expressed by the second term in Eq. 3.11. Thus, I (implicitly) take $\mathrm{Var}_N[I_i] = 0$, thereby also eliminating the third term. This situation models a homogeneous population in terms of models of the world, such as for individuals of the same species in a common habitat, or workers in the same industry, with similar training. I will return to the more general case and discuss future opportunities in the discussion at the end of the chapter.

For a population of agents independently sampling a shared multinomial environment, the initial variance in growth rates is given by the variance in the initial binomial parameter, σ_x^2. The dynamics of the binomial parameter variance for a population of size N is (SM 22)

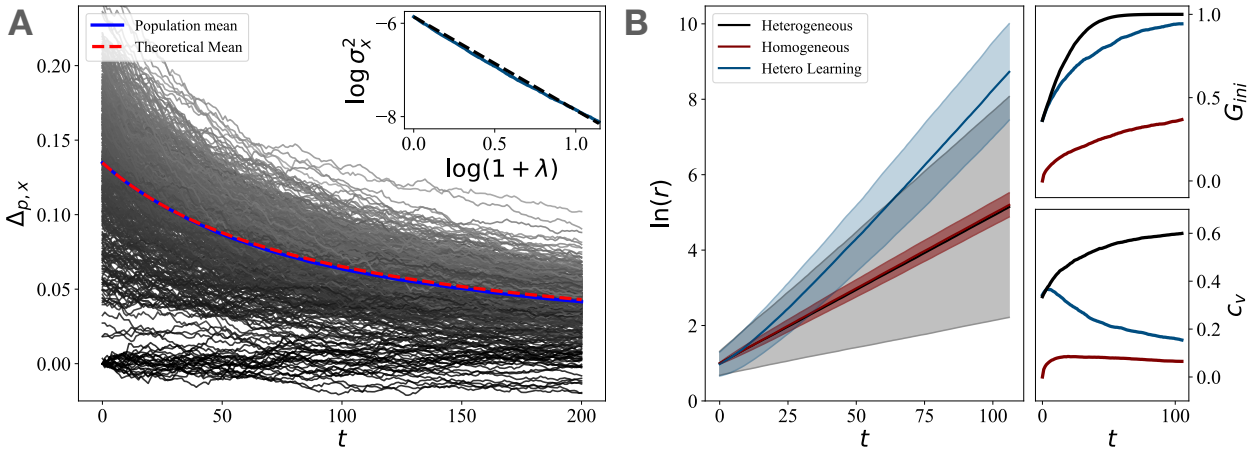

Figure 3.4: Monte Carlo simulations of a population undergoing growth dynamics, with and without Bayesian learning, with parameters of mean growth rate, $\bar{\gamma} = .04$, and standard deviation, $\sigma_\gamma = .641\bar{\gamma}$. **A**. The simulated and theoretical means of x converge to p, thus maximizing growth rates. The parametric variance, σ_x^2, (blue) follows the theoretical prediction (red). The linear behavior log-log plot demonstrates the power law behavior of σ_x^2. **B.** The mean resources of three population types are plotted with a shaded region providing 95% CI bounds for single agent trajectories. Heterogeneity broadens the range of possible wealth values, while learning increases both mean growth while narrowing the shaded region relative to no inference. Agent learning slows the increase in the Gini coefficient introduced by heterogeneity and reduces the coefficient of variation.

$$\mathrm{Var}_N\big[x_i(t)\big] \equiv \Big\langle \big[x_i(t) - \langle x(t)\rangle\big]^2 \Big\rangle = \frac{\sigma_x^2}{(1+t/kl)^2}, \tag{3.12}$$

where $\langle x(t)\rangle = \frac{1}{N}\sum_i x_i(t)$. Assuming a population of entirely conservative (or aggressive) agents, such that all growth rates map to a unique binomial parameter, we can approximate the variance in growth rates, $\sigma_\gamma^2(t) = \Big\langle \big(\gamma[x_i(t)] - \gamma[\langle x(t)\rangle]\big)^2 \Big\rangle$, by Taylor expanding the second moment of the resource distribution. The approximation carried out in SM 38 shows that the growth rate variance decreases asymptotically in polynomial t^{-2} time [177]. Figure 3.4.A. demonstrates that in a population of agents sampled from a Gaussian distribution of growth rates and resources learning their environment, $\Delta_{p,x} = p - x(t) \to 0$ as $t \to \infty$, and individual binomial parameters converge to the optimal value. At the population level, there is an agreement between the empirical population mean and theoretical mean trajectory, calculated by evolving $\langle x(t)\rangle$ using Eqn 3.10. Similarly, the empirical population variance in

x matches the theoretical power law prediction given by Eq. 3.12.

These results show that learning a shared, stationary environment reduces growth rate variance on the same time scale as the dynamical effects introduced by growth rate variance [140]. This shows that fast learning (sufficiently low k) equalizes information access, and is a suitable mechanism for reversing the long-term effects of heterogeneous growth on inequality.

I demonstrate these features of the dynamics by comparing the statistics of resources across Monte Carlo-simulated populations. I first use homogeneous initial conditions, then heterogeneous initial conditions with and without inference. To measure the increase in inequality, I track the Gini coefficient, denoted G_{ini}, which varies between zero –for uniformly distributed resources– to 1, for maximally unequal wealth distributions. (For a lognormal distribution, such as in the GBM model, $G_{ini}(t) \approx \mathrm{Erf}[\sigma_r^2(t)]$). Additionally, I measure the relative increase in standard variation to the mean of resources via the coefficient of variation, $c_v = \sigma_r/\langle r \rangle$. More on this analysis is given in [140].

The resource time evolution shown in Figure 3.4B demonstrates that growth rate heterogeneity dramatically broadens the wealth distribution, in agreement with [140]. Accordingly, heterogeneity increases G_{ini} and c_v as compared to a homogeneous population. The introduction of learning increases the average growth rate in a heterogeneous population, as demonstrated by the higher mean wealth, while reducing the variance in resources. The former slows the rapid increase G_{ini}, while the combination of both reduces c_v to levels comparable to the homogeneous trajectory, confirming that learning reverses the effects of heterogeneity on inequality.

While this simplified model does not capture the nuanced effects of educational systems or skill heterogeneities implied in real societies, the connection between convergent learning in a population and growth is general and provides a sound theoretical basis for the observed benefit of education on national growth, human capital, and inequality reduction [117, 148,

176].

3.4 Discussion

In this chapter, I developed a statistical dynamical theory for the origin of resource growth rates in populations of learning agents experiencing a shared stochastic environment. I showed that an agent's growth rate is, in the limit of many decisions, the quantity of mutual information between their signal and the environment and that learning through Bayesian inference provides a natural and necessary (optimal) mechanism for increasing agents' growth rates, managing volatility, and reducing growth disparities across populations over time. I demonstrated that in the particular static case (without learning), this framework reproduces GBM models widely used in wealth dynamics and inequality studies, and provides models for their parameters. When agents can learn, their parameters become optimal over time and acquire formal interpretations in terms of information.

The present treatment answers an important open question established in the previous chapter on how to mechanistically control variances in growth rates across a society while maximizing learning and growth and generally enriches the typical modeling schema of wealth dynamics by incorporating agents' subjective choices in a structured, stochastic but knowable environment. Towards the end of this dissertation, this work adds to the foundations necessary for incorporating formal models of information and adaptive agent behavior in statistical mechanics, establishing a microfoundation for broadly applied nonequilibrium statistical mechanical model.

There are a number of interesting developments that this theoretical framework suggests for modeling more realistic, particular situations. First, learning is never quite uniform across populations or time, varying across the life course, with some agents being able to dedicate more time and effort to it than others. This issue can be modeled by making inference rates dynamic and heterogeneous, for example, through coupling to agents' socioeconomic

status (SES) or age. Importantly, lower SES has been shown to be correlated with the presence of stressors that inhibit the cognitive ability of people to learn [75, 112, 243], while higher SES correlates with better educational outcomes [18, 40, 163]. Coupling learning rates with SES would alter the population's learning trajectory and potentially attenuate its effectiveness in reducing information and wealth inequality. Moreover, my analysis has assumed that each agent samples identically distributed signals. In reality, people across different structural positions in social networks, for example, associate with place, gender, or race/ethnicity, typically have differential access to signals (opportunities), with implications for what they can learn and for resulting social equity. Along with different signals, different agents may have different models of the world, which are naturally incorporated in the scheme developed here by different likelihood functions in Bayesian learning. I have shown in general how such heterogeneities among agents will result in inequalities in their growth rates, but many interesting situations remain to be explored in the future. Finally, real societies also feature interactions between agents and planners that redistribute wealth, economic rents, and heterogeneous frictions that further shape wealth dynamics. The work developed here emphasizes the importance of understanding these policy choices and socioeconomic phenomena through the lens of how they affect specific wealth dynamical parameters, namely initial wealth statistics versus growth, including average growth rates and volatilities across populations. Future studies of the origins of inequality and social equity should consider these structural complexities from the general point of view of access to information and learning, and the specific analytical tools that they introduce.

Second, from the point of view of maximizing resources, there are familiar trade-offs between learning and investing. These can be modeled in terms of the inference process divided into passive experiential learning, resembling the "learning by doing" featured above, and, additionally, emulating formal, institutional education wherein agents sacrifice short-term wages to more rapidly acquire information. These considerations define agent trade-offs

between actively exploring and passively exploiting the environment, an important research topic in both experimental neuroscience and machine learning [145, 236]. Furthermore, while information is a non-rival quantity that can be made available to a society with minimal cost of sharing or degradation from use, the generation and dissemination of information through teaching is a costly process that can produce additional non-trivial dynamics. Agents must also consider the cost and benefits of seeking education in non-stationary environments, where the value of information may fluctuate or decay over time. Expressing the social costs of education through mechanisms of finite learning resources could help explore trade-offs in investing in human capital over various timescales of learning and environmental evolution [189, 218] and help determine when they are worth it - for individual agents and societies - in inter-temporal settings.

Third, tracking individual agent dynamics under constraints of finite (varying) lifespans can help determine the effects of generational wealth transfers on inequality, and provide insight into life-course strategies [74] and issues of valuing (and discounting) the future. Thus, an extended framework can help us explore the scope of education under the discounting of delayed resources by longevity and lived volatility [116]; including the implications of costs and expected earnings with or without an education over time. This research topic is discussed in the final chapter of this dissertation as an active field of research.

Lastly, agents in this model experience the same environment and learn the same information, whereas actual communities specialize in different, complementary skills that may minimize knowledge redundancy. These information complementarities and exchanges are known commonly in the social and ecological sciences in terms of the division of labor and knowledge [24]. How agents decide which information to learn and what profession to choose based on their environments begets different growth rates across a population, altering emerging inequality and influencing how social groups cooperate or compete across community or institutional social levels [87]. Cooperation among agents with synergistic information in a

stochastic environment has been shown to produce non-linear additive effects on aggregate information [25], suggesting that cooperative agents would experience larger growth rates when coordinated, compared to the sum of agents acting independently [138, 201]. In the next chapter, I study this connection between social behavior and growth from the point of view of information and learning, and provide insight into the origins of sharing resources and information in stochastic growth settings.

CHAPTER 4

COLLECTIVE ACTION AND GROUP FORMATION

In the prior chapter, I established how information underpins the optimal decisions agents make in noisy environments, and the growth they experience. In this chapter, I will demonstrate how this theoretical insight can improve our understanding of the quantitative benefits to collective behavior. I begin by motivating a theory for cooperation among heterogeneous agents by reviewing progress and shortcomings in the organizational science, ecology, and game theory communities. I then modify the theory introduced in the prior chapter to support collective decision-making. Towards the end of the chapter, I will demonstrate its effects with a simple model and explore its implications [1].

4.1 Introduction

Collective behavior is a general feature of biological and social systems. It mediates the survival and evolution of populations under resource constraints, competition, or predation in natural systems [190] and the formation and persistence of social organizations in human societies [229]. Much past work has modeled collective dynamics using homogeneous interaction rules, common to all agents, which are often phenomenological. While these models have produced diverse insights, they typically lack a theoretical foundation to explain how specific social behavior emerges among individual agents with heterogeneous information and behavior. Thus, significant knowledge gaps remain in most realistic situations, where agents with distinct but potentially complementary traits act collectively to maximize their joint growth (fitness, wealth) in knowable but stochastic environments.

Some examples help illustrate the present situation. Game theorists and ecologists have considered many different cooperative interaction schemes [208] and explored evolutionary

1. This chapter is adapted from Kemp, J. T., Kline, A. G., & Bettencourt, L. M. (2024). Information synergy maximizes the growth rate of heterogeneous groups. *PNAS nexus, pgae072*

stable behavior [114], particularly on networks [129, 193, 220], where optimal behavior is identifiable under given interaction rules. Elaborating these schemes by introducing higher order interactions has broadened our understanding of more complex social networks [6, 12, 46, 103], and their dynamical phase-stability under varying interaction strengths [81]. Researchers have also studied, both theoretically and in the laboratory, how memory of previous interactions influences agents' preferences for future encounters [106, 107, 168, 207], the spread of social crises across distance [155], and the formation and scaling properties of social collectives [44, 203], such as cities [26, 216].

In addition to interaction rules and associated payoffs, collective dynamics is predicated on maximum principles, which specify agents' preferences in view of a goal and thus render their behavior intelligent (optimal). For example, inclusive fitness theory, which assumes a reproductive benefit to cooperation because of shared genes [113, 191] has been studied in mixing populations and over networks [183] where it predicts population benefits to cooperation through several forms of reciprocity [201]. More recently, researchers have studied resource pooling in models of growth as a means to minimize environmental uncertainty and associated loss of fitness among agents experiencing independent fluctuations with shared statistics [160, 196]. Such approaches remain limited by the association between collective behavior and (genetic) homophily. Still, they can help explain the existence of phase transitions in cooperation networks [44, 81], and specify agents' plausible behavioral patterns [106], even if doubts remain about inclusive fitness's predictive power [181].

Generally, however, most current quantitative frameworks fail to address collective dynamics when agents remain heterogeneous across skills, knowledge, and behavior [83, 192, 228]. Developing more general approaches to collective behavior that include adaptation along with heterogeneity is a crucial step towards understanding how agents self-organize in more complex and dynamical environments, where specialization and the division of labor and knowledge become key.

Adaptive behavior requires agents to acquire and process information over time [69, 86] in response to their environments and each other. In realistic situations, limited experience, specialization costs, and physical limitations of effort, energy, and time, all prevent agents from perfecting their knowledge of complex environments [174, 233]. A natural way to mitigate these individual limitations is to pool knowledge across agents leading to the formation of social organizations [126], and the division and coordination of labor in terms of their behavior [55]. This is widely observed in human organizations and animal social behavior starting with the division of labor by age and sex.

By working jointly to predict characteristics of their environment [228] and gather resources, groups of agents can maximize their collective fitness even when each individual has very limited knowledge. In a setting where there are resource returns to successful prediction and behavior, information of the state of a statistical environment determines the fitness of the population [28, 141], though there are questions about how such benefits emerge quantitatively [25]. Here I formalize the calculation of these social benefits in terms of the properties of information and show how maximizing knowledge complementarities (synergy) maximizes the long-term growth rate of collectives. Specifically, I derive an expression for the additional payoff to cooperative behavior in terms of the joint information synergy about the agents' dynamical environment.

The aggregation of dispersed, tacit information among a group of agents has long been proposed as the principal role of economic markets [118], operating through the price mechanism. In such settings, a public price forms as the result of the allocations of traders with diverse knowledge, buying and selling an asset according to their beliefs (estimates) of its value. Several types of markets, both centralized [108, 173, 244] and decentralized [34, 70, 246] have been discussed as efficient aggregators of information in this sense, but fundamental objections have also been raised [109]. Information, in the sense of this *efficient markets hypothesis*, usually reflects only average beliefs among traders [245]. By contrast,

my approach shows how dispersed knowledge can be combined in optimally predictive ways.

These results lead us to introduce the principle of maximum synergy, which maps the maximization of pooled resource growth rates into optimal social interaction structures. This work adds new dimensions to the study of collective dynamics by connecting the structure of groups to that of information in complex environments mediated by agents' diverse subjective characteristics, such as their present knowledge and information acquired as the result of diverse experiences throughout their life course.

4.2 Theory of Collective Growth

I start by demonstrating how the benefits of collective action emerge from pooling information in synergistic situations. Synergy means the combination of behavior, knowledge, and skills that complement each other toward a goal. This concept is necessary for creating effective organizations that embody complex information [228], but it is often not sufficiently formalized in common language, such as in discussions of innovation [94] or firm structure.

Here, I will refer to synergy as an explicit information-theoretic quantity that measures the additional predictive power that a group acquires upon pooling its agents' information, relative to the knowledge of each individual separately. This quantity has been introduced some time ago in the context of studying circuits in information processing systems [31, 217], and has provided a framework for studying higher-order neuron interactions in the brain [239], and causality and information in complex systems [169, 238]. As I will show, synergy results formally from the conditional dependence between the probability of predictive signals distributed in a population and events in a shared environment. The gain in predictive power from agents pooling information as collectives allows them to obtain additional resources from a knowable environment beyond what agents alone can do, thus boosting their fitness or productivity.

It follows that collectives that seek to maximize their resources over long times must

combine the information from their agents' individual models of the world in a way that accesses the most synergy. Groups that do not know *a priori* how to realize their synergies must discover how to do so, by adjusting their collective knowledge and interaction structure while observing outcomes of their environment in an iterative learning process. After developing the general framework for group formation and collective growth across group sizes, I demonstrate a model environment that exhibits synergy using logic gates. I will also demonstrate how synergy scales with the number of unique signals in a collective, and how specific combinations of signals affect the average growth of resources for the group.

4.2.1 Collective Growth in Synergistic Environments

Consider a population of N agents, each with initial resources $r_i,\ i = 1, \ldots, N$ that can be (re)invested into the set of outcomes of their environment to generate returns. Each agent has access to a private signal (their knowledge), $s \in S$, which is used to predict the state of the environment and make resource allocations to possible outcomes $e \in E$. This signal may represent several different processes such as sensory input, or a lead retrieved from memory. With optimal parameterization of a model of the environment, $P(E|S)$, an agent's optimal investment strategy leads to an average resource growth rate (over time) $\gamma = I(E;S)$ [141], where $I(E;S)$ is the mutual information between environmental states E and the agent's signals S. (I am working in units of units time $t = 1$, for simplicity.) Agents with better models (and better statistical estimations of $P(E|S)$) thus experience higher average growth rates.

Now, define the agent's environment more explicitly, by a set of l distinct signals with unique statistics, $\boldsymbol{S} \equiv \{S_1, \ldots S_l\}$ as $P(E|\boldsymbol{S})$, with marginals of events $P(E)$ and signals $P(\boldsymbol{S})$. The joint information that the universe of signal, $\boldsymbol{S}$, has on E is at least equal to each of the signals S_j, that is $I(\boldsymbol{S};E) \geq I(S_j;E)$, for all j. Generally, this inequality is strict if the conditional information $I(\boldsymbol{S}|E) > I(\boldsymbol{S})$ [31, 217]. Compute the total information by

summing over the mutual information between each of the signals independently, subtracted by an interaction term across them,

$$I(E;\boldsymbol{S}) = \sum_j I(E;S_j) - R_P. \tag{4.1}$$

The coefficient of redundancy, R_P, measures the strength of this conditional dependence across larger sets of signals (two, three, etc). It is defined in App. A

$$\begin{aligned} R_P = &\sum_{j>k=1}^{l} R(E;S_j;S_k) + \sum_{j>k>m=1}^{l} R(E;S_j;S_k;S_m) \\ &+ \cdots + R(E;S_1;\ldots;S_l). \end{aligned} \tag{4.2}$$

The coefficient of redundancy can have a positive or negative value, indicating different conditional relationships between the signals and environmental states. When $R_P > 0$, there is information between signals irrespective of environmental events. This means that signals are partially *redundant*, and consequently, there are diminished returns to pooling information as $I(E;\boldsymbol{S}) < \sum_j I(E;S_j)$. Conversely, when $R(E;\boldsymbol{S}) = 0$, the signals are statistically independent, and the benefits of pooling information increase linearly with the information of each signal on the environment but there is no synergy. Finally, when $R(E;\boldsymbol{S}) < 0$, there is conditional dependence of the signals on the environment. This is called *synergy* and yields a superlinear benefit to pooling information in the number of agents, above and beyond the information contributed from each signal individually.

Group formation and collective decision-making

I have now defined individual resource growth rates as a quantity of information and discussed how information can be aggregated across different signals to express their synergy relative to states of the environment. Now we can explore how agents with different signals can pool

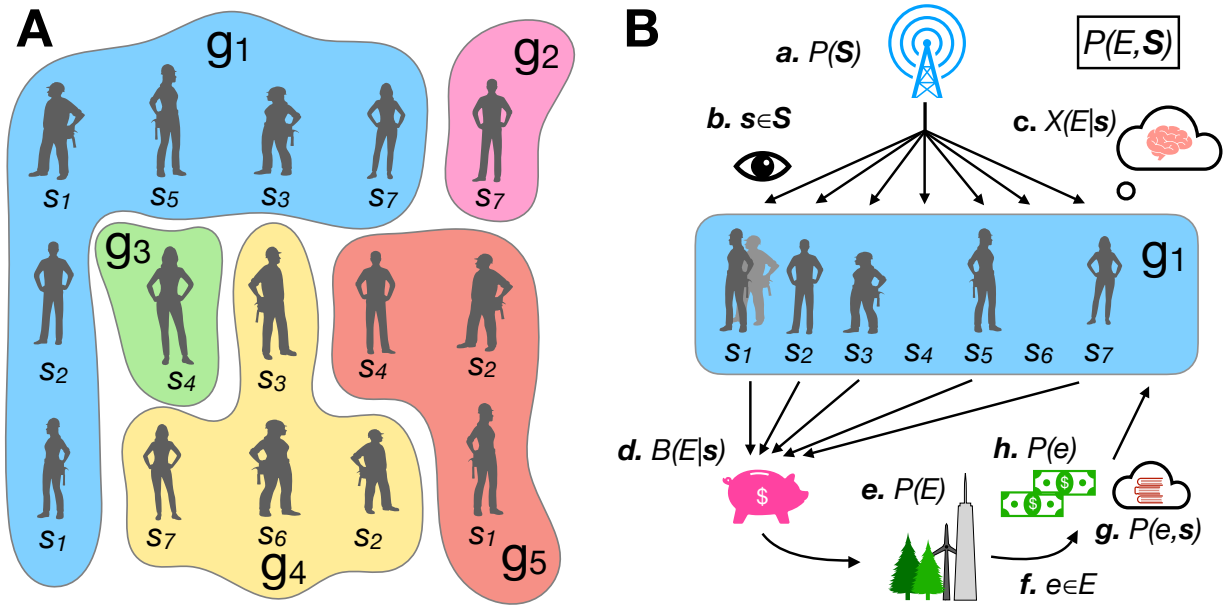

Figure 4.1: Groups of agents with different signals grow resources based on the information between their signals and states of the environment. **A.** Groups, denoted g, are composed of an arbitrary number of agents. Each agent belongs to only one group and can observe and contribute one signal to the group. A group contains k_g unique signals. **B.** At each time step, (a) the group's private channel outputs a signal $\boldsymbol{s} \in \boldsymbol{S}$ with probability $P(\boldsymbol{s})$. (b) Each member of the group observes their signal s_j and (c) the group consults their collective belief for the conditional outcome probability of the environment, $X(E|\boldsymbol{s})$. (d) The agents make proportional resource allocations on all possible outcomes $B(E|\boldsymbol{s})$. (f and e). The true event $e \in E$ is observed in the environment with probability $P(e)$, and (g) the agents receive payouts proportional to the marginal probability of e.

information together as coordinated groups, and access the synergy in their environment through collective decision-making.

Consider the undirected hypergraph $H = (A, G)$ of vertices, A, and hyperedges G. Consider a discrete number of vertices, $A = \{a_1, a_2, \ldots, a_N\}$, where a_i identifies agent i. The set of hyperedges, $g \in G = \{1, 2, \ldots\}$, called groups, defines the number of cooperating collectives. A hyperedge connects $1 \leq N_g \leq N$ agents. Assume that agents can only belong to a single group. Therefore, by construction, $\sum_g N_g = N$ and the sum over all nodes of every hyperedge yields the number of agents in the population. There exist two extremes of cooperation. First, when a single hyperedge spans every node, meaning all agents pool information in a single group. In the limit of no cooperation, $N_g = 1$ for all g, and no agents pool information. In this case, the dynamics of the model are similar to previous work [141].

Let $\boldsymbol{S}_g$ be the set of unique signals held by the agents of a group g to be pooled, such that $\boldsymbol{S}_g \subseteq \boldsymbol{S}$. The number of *cooperants* is defined by the number of unique signals, $|\boldsymbol{S}_g| = k_g$, and is bounded by $1 \leq k_g \leq l$. When $k_g = l$ and the group has a complete signal, the collective can make maximally informed decisions. Conversely, when $k_g < l$, the signal is considered *incomplete*, and the collective can only interpret and act on a subset of signals. As we will see, the number of unique signals a collective can observe determines the amount of information they can access.

Now that agents are organized into groups of various sizes, we can discuss how agents pool their information to make collective decisions and grow their resources in dynamic environments. At every time step, a collective with access to all signal types observes a unique private signal $\boldsymbol{s} = \{s_1, \ldots, s_l\} \in \boldsymbol{S}$. Each agent then allocates its resources r_i on events according to collective g's allocation matrix $B(E|\boldsymbol{s})$. As the event e is observed, the agent is rewarded with returns w_e to the fraction of resources invested in e, $B(e|\boldsymbol{s})$. In the limit of many sequential investments n, the average growth rate of resources converges to

$$\gamma = \frac{1}{n} \log \frac{r_n}{r_i} \approx \sum_{e,\boldsymbol{s}} P(e, \boldsymbol{s}) \log \big[B(e|\boldsymbol{s}) w_e \big]. \tag{4.3}$$

The optimal investment in the large n limit is the conditional probability of the event given the signals, $B(e|\boldsymbol{s}) = P(e|\boldsymbol{s})$. When the rewards are "fair", and $w_e = 1/P(e)$, the optimal growth rate is given by the mutual information [137] defined in equation 4.1, $\gamma = I(E; \boldsymbol{S})$.

The typical collective may not have a complete signal, and instead may only observe and interpret a subset of all unique signals $\boldsymbol{S}_g$. Their optimal allocation, given by $P(E|\boldsymbol{S}_g)$, then has mutual information $I(E; \boldsymbol{S}_g) \leq I(E; \boldsymbol{S})$, with equality only if the omitted signals are completely redundant with present signals. Unless there are redundant signals, an incomplete group is guaranteed to have suboptimal information and growth rate.

Agents *a priori* may also not have perfect knowledge and must invest using their best estimate of the true conditional probability, $X(E|\boldsymbol{S}_g) \neq P(E|\boldsymbol{S}_g)$. In this case, the collective's average growth will be submaximal by the number of signals and lack of information on signals and is described by

$$\gamma_g = I(E; \boldsymbol{S}_g) - \mathrm{E}_{\boldsymbol{s}_g}\big(D_{KL}\big[P(E|\boldsymbol{s}_g) || X(E|\boldsymbol{s}_g) \big] \big), \tag{4.4}$$

where $\mathrm{E}_{\boldsymbol{s}_g}$ is the expectation value over the states of the group's signals, and $D_{KL}\big[P(E; \boldsymbol{s}_g) || X(E; \boldsymbol{s}_g)\big] = \sum_e P(e|\boldsymbol{s}_g) \log\big(P(e|\boldsymbol{s}_g)/X(e|\boldsymbol{s}_g)\big) \geq 0$ is the Kullback-Leibler divergence, an information measure expressing how similar the distributions are. This result shows that collectives with both a better model as reflected by the first term, a better characterization of the model and its various synergies by the second, and a more complete signal, will experience higher growth rates. Furthermore, $\gamma_g < \gamma$ unless g is the full set of signals, so it is typically valuable to add more signals to the group. This setup is illustrated in Figure 4.1.

4.2.2 Maximum synergy principle and optimal growth

These results introduce important considerations for how collective innovation and growth determine strategies for group formation. In theories of cooperation such as kin selection [72] and scalar stress [132], group formation is advantaged by member relatedness and disadvantaged by unfamiliarity. This is intuitive in many situations, as agents are more likely to cooperate when they are more certain others will reciprocate [180], and cooperating with similar agents may naturally minimize this uncertainty. Equation 4.4 counters this intuition by defining an explicit benefit to cooperating with dissimilar agents across heterogeneous, complementary skills and information. Specifically, a group with more synergistic signals, as defined through the conditional dependence of their decisions on states of the environment, will experience higher growth. So, even if there are additional coordination costs for more heterogeneous agents, there is now a possibility that cooperation will emerge as there are also greater informational benefits, formalizing intuitive ideas about the value of diversity [186].

The beneficial contribution of synergy to the growth rate of resources provides an important input to models of random multiplicative growth, such as those commonly used to study wealth dynamics and mathematical finance. In its simplest form, the stochastic growth rate in such models is characterized by its first two temporal moments. The average over time, η, and the resource temporal standard deviation (volatility), σ, combine under Itô integration to give the actual growth rate $\gamma = \eta - \sigma^2/2$. Maximizing this growth rate (as a positive quantity) entails maximizing η and minimizing σ, which at the individual agent level can be achieved by (Bayesian) learning over time [141].

At the population level, it has been proposed that pooling resources in groups would naturally emerge as a means to reduce σ, when growth rate fluctuations are independent across agents, and thus maximize γ [78, 196].

Our results introduce a different possibility of cooperation, through pooling information in structured groups, that maximizes η (and γ) through synergy effects. Thus, to maximize

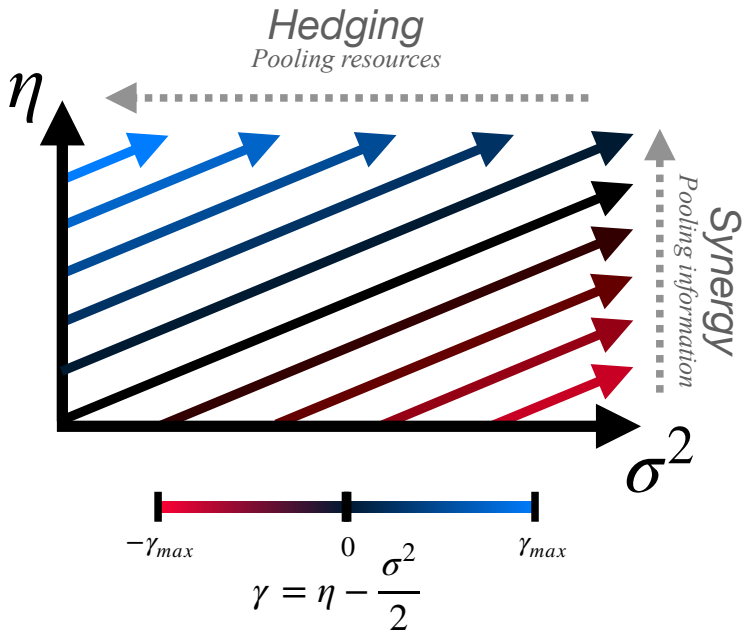

Figure 4.2: Complementary strategies for increasing the long-term growth rate of resources from the environment in stochastic growth models. Pooling resources can reduce volatility through a hedging strategy while pooling information creates synergy to increase average growth rates. The lines represent contours of constant average growth rates γ.

γ, agents should pool information with the most diverse set of collaborators possible to access the most mutual synergy viz. the environment. This *maximum synergy principle* defines the benefit of intelligent collective behavior in complex environments where there are agent-level limitations to knowing the environment fully and where mechanisms of the division of labor and knowledge are favored. This principle is general and applies across levels of cooperation, whether it be individuals matching skills to form groups, or specialized groups organizing into more complex collectives [25], all the way to large-scale societies.

Generally, these two strategies, information synergy versus resource pooling under independence, are distinct modes of cooperation over which groups can maximize γ, as demonstrated in Figure 4.2.

As we will see later, the decision of whom to cooperate with is not trivial, as different combinations of signals may yield varying synergies. This means that under constraints to group size such as from cooperation costs per connection, groups satisfying the maximum

synergy principle must intelligently select which signals and agents to integrate, and which to exclude as redundant.

Furthermore, collectives may not *a priori* know the optimal allocation strategy that leverages the synergy available to their signals, meaning that intelligent collective behavior must itself be learned over time and by exploring the best possible matchings. I will now develop the dynamics of how a group maximizes its synergy given a set of signals.

Synergy maximization through Bayesian inference

Bayesian learning is the optimal strategy to incorporate new information from observed events into the estimate of conditional probabilities, such as those of environmental states given agents' signals [141]. Agents can also learn the synergy embedded in their environment in groups by collectively weighing their conditional observations across their individual signals. A group wanting to maximize their synergy must then update their conditional relationship through a Bayesian inference process

$$X_n(e|\boldsymbol{s}) = AP(\boldsymbol{s}_n|e_n)X(e_n) = \left[\prod_{i=1}^{n} \frac{P(\boldsymbol{s}_i|e_i)}{P(\boldsymbol{s}_i)}\right] X(e), \tag{4.5}$$

where the normalization $A = (\int de_n P(\boldsymbol{s}_n|e_n)X(e_n))^{-1}$. Take the prior probability, $X(e_1) = X(e)$, because I am assuming that the environment is stationary or at least slowly changing relative to groups' learning rates.

Bayesian inference converges $X(E|\boldsymbol{S}) \to P(E|\boldsymbol{S})$ over time, decreasing the information divergence, and maximizing synergy and average growth. For groups with incomplete signals, the information acquired through learning is still bounded by what is available in the incomplete signal space.

I have thus far defined collective growth in terms of information synergy, and shown how agents can learn as a collective to increase their growth rate over time. I will now illustrate

these general results using a model based on logic circuits.

4.3 Modeling Synergy with Logic Circuits

Logic circuits have been used extensively as models for synergistic interactions [25, 31, 217]. This is because their outputs are predicted by combinations of inputs, much like events are predicted by combinations of signals. Among other logic circuits (like AND or OR), the XOR gate is unique in that information between inputs and outputs only exists as synergy across all inputs [130]; no individual input has mutual information with the output.

In the following section, I will show how modifying the XOR gate relaxes this condition, such that information exists for any input and scales on average with the number of cooperating signals. Similar to [141], while this model will be used to study synergy in a simplified setting, the theory is defined for general dynamical environments.

4.3.1 The Uniform XOR Gate

Consider the space of statistically independent binary signals $s_j \in 0, 1$, such that a sample set $\boldsymbol{s}$ has uniform probability $P(\boldsymbol{s}) = 2^{-l}$. Assign each input $\boldsymbol{s}$ a binary event, $e \in 0, 1$, using the generalized XOR rule, $e = M_2(\boldsymbol{s}) \equiv \left[\sum_{j=1}^{l} s_j\right] (\text{mod } 2)$ with binomial probability $p_{\boldsymbol{s}}$. From the sets of sampled signals, $\boldsymbol{s}$, and binomial coefficients $\boldsymbol{p} = \{p_{\boldsymbol{s}}\}$, I define this generalized XOR circuit as a joint distribution on signals and events as

$$P(E, \boldsymbol{S}|\boldsymbol{p}) \equiv f(\boldsymbol{p}, l) = \frac{1}{2^l} \prod_{\boldsymbol{s}} (p_{\boldsymbol{s}})^{M_2(\boldsymbol{s})} (1 - p_{\boldsymbol{s}})^{1 - M_2(\boldsymbol{s})}. \tag{4.6}$$

This distribution is called the uniform XOR (UXOR). It performs a unique, l dimensional XOR gate on each input $\boldsymbol{s}$ with probability $p_{\boldsymbol{s}}$. When $p_{\boldsymbol{s}} = 1$ for all input permutations, this circuit behaves deterministically like an XOR gate, and the complete group has 1 bit

of information. In the limit of $p_{\boldsymbol{s}} = .5$, this no longer models a logic gate as the output is uncorrelated to the inputs. The truth table of this circuit is shown in Figure 4.3A for an environment with two signals.

Information scaling in the UXOR environment

With this explicit choice of distribution, this section explores the information quantities that will define a group's growth process. For simplicity, I choose a uniform prior for the distribution of $\boldsymbol{p}$, but in principle any prior distribution is admissible. The information available in the environment measures the maximum average growth rate a group with a complete signal can experience. When averaged over all configurations of $\boldsymbol{p}$, the information is given by $I(E; \boldsymbol{S}) = \log 2 - 1/2 \approx .28$ bits (C.2.1).

For groups with incomplete signals (when $k_g < l$), I compute the information by marginalizing equation 4.6 over the $\lambda_g = l - k_g$ signals unavailable to the group. The procedure for marginalization is defined in App. C, but in general, marginalization of one signal halves the size of the parameter space $\boldsymbol{p}$ that describes the distribution. The average information for an incomplete signal is approximately [143]

$$I(E; \boldsymbol{S}_g | \boldsymbol{p}) \approx 2^{-\lambda_g} \left(\log 2 - \frac{1}{2} \right). \tag{4.7}$$

Average information scales exponentially, $\sim 2^k$, as more signals are included. The mutual information of the complete signal is independent of the number of signals, so the information of a single signal must converge to zero in the limit of large l.

The exponential scaling of the information with the number of cooperants is demonstrated in Figure 4.3B, as lines on a logarithmic scale for environments of increasing l. The curves are computed by Monte Carlo sampling circuits for l signals by measuring the information after $\lambda = l - k$ marginalizations.

A

s_1	s_2	$P(e{=}0)$	$P(e{=}1)$
0	0	p_1	1-p_1
0	1	1-p_2	p_2
1	0	1-p_3	p_3
1	1	p_4	1-p_4

B

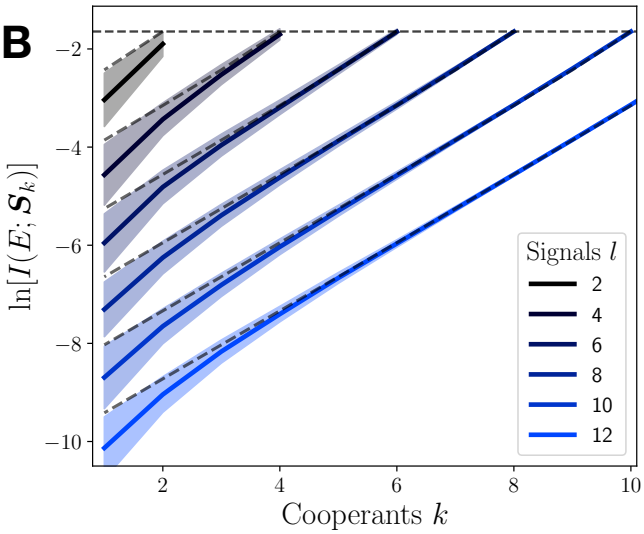

Figure 4.3: The UXOR model provides an environment for exploring synergy across groups of arbitrary size. **A.** The UXOR circuit, demonstrated by the modified XOR symbol, and its truth table for $l = 2$. **B.** The information of a circuit of size l scales exponentially in cooperants, k.

Growth and group learning

Until now I have explored the mean behavior of this environment subject to a uniform prior. In general, collectives do not have perfect information on a single prior. In this case, their inaccurate guess for the set of binomial coefficients is parameterized by $\boldsymbol{x}_g \equiv \{x_{\boldsymbol{s}_g}\}$, indexed by the signals available to the group $\boldsymbol{s}_g \in \boldsymbol{S}_g$, and the collective's likelihood model becomes $X(e|\boldsymbol{s}_g) = f(\boldsymbol{x}_g, k_g)$. The information divergence term of equation 4.4 becomes the divergence between $f(\boldsymbol{x}_g, k_g)$ and $f(\boldsymbol{p}_g, k_g)$, where $\boldsymbol{p}$ has been projected into the subspace spanned by $\boldsymbol{S}_g$, averaged over all signals $\mathrm{E}_{\boldsymbol{s}_g}[D_{KL}] = \langle p_{\boldsymbol{s}_g} \log\big(p_{\boldsymbol{s}_g}/x_{\boldsymbol{s}_g}\big) + (1 - p_{\boldsymbol{s}_g}) \log[(1 - p_{\boldsymbol{s}_g}/(1 - x_{\boldsymbol{s}_g})]\rangle$ here angle brackets denote sample averages over the binomial values. Subtracting the mutual information by this term yields the growth rate under imperfect, incomplete group information.

$$\gamma_g = \big\langle p_{\boldsymbol{s}_g} \log x_{\boldsymbol{s}_g} + (1 - p_{\boldsymbol{s}_g}) \log\big(1 - x_{\boldsymbol{s}_g}\big)\big\rangle + \log 2, \tag{4.8}$$

I have so far described growth rate dynamics under a stationary $\boldsymbol{x}_g$. To illustrate growth dynamics under group learning, I turn to the Latent Dirichlet Allocation (LDA) model. Through a categorical description of pairs of events and signals, agents experience average dynamics to $\boldsymbol{x}_g$ in the limit of high sampling rate $\omega = n/t \gg 1$

$$\boldsymbol{x}_g(t) = \frac{\boldsymbol{p}_g t/2\kappa + \boldsymbol{x}_g}{1 + t/2\kappa}, \tag{4.9}$$

where κ defines the Bayesian update time. The details of LDA are given in Ref. [141] and provide parametric dynamics that converge to full information as a power law in time, in stationary environments.

To study resource dynamics in the UXOR environment, I simulated agent investments in a Monte Carlo sampled environment. I randomly assigned $N = 5000$ agents signals in an $l = 4$ environment, then randomly assigned them to groups sized $l \leq N_g \leq 11$. This results in an ensemble of groups with cooperants $1 \leq k_g \leq 4$. I reveal Bernoulli-sampled signals to the groups, whose agents collectively decide on which events to allocate resources to. For each group, I track the resources of a representative agent, informed by the group, investing their individual resources through time. The full details of the setup are provided in the supplementary materials.

Figure 4.4 illustrates the results of this simulation. In subfigures A and B, the Monte Carlo simulated means are shown as solid lines, with 95% Confidence Interval (CI) shaded regions. Theoretical means are computed from the initial population configuration using equation 4.9, plotted as dashed lines, with hash-filled uncertainty regions. Simulated groups have randomly assigned members with uniformly assigned signals, where $N = 2000$. The more unique signals a group can access, the more they can learn, and the more resources they acquire over time. A high signal-to-noise ratio when $k_g = 1, 2$ causes growth rates to be lower than the theoretical mean, and cumulatively results in fewer resources over time.

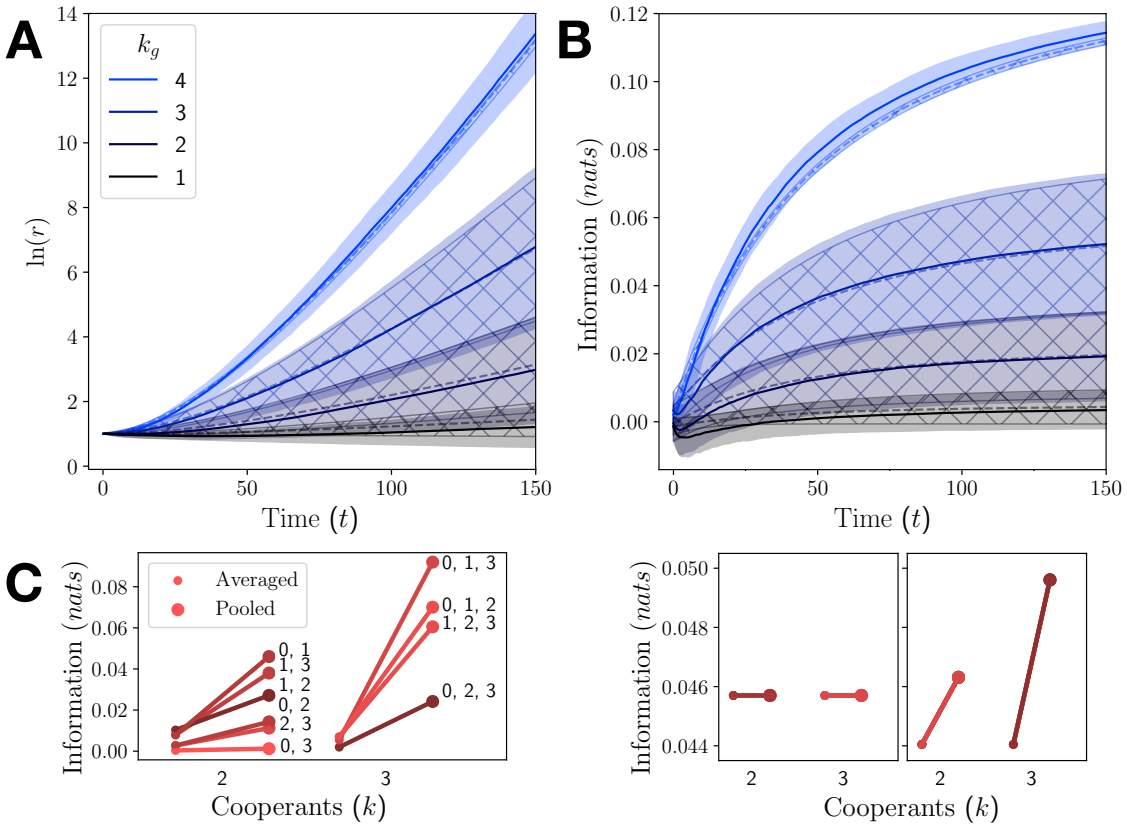

Figure 4.4: Groups learning an $l = 4$ environment using more unique signals acquire more resources and information, but combinations of signals have unique amounts of information. **A**. Temporal resource trajectories, grouped by number of unique signals in the corresponding group show that growth increases with the number of signals. **B**. Groups with more signals can gather more information from the environment. There is high variability when $k_g < l$, as different combinations of signals access different amounts of information. **C**. *Top* For $k_g = 2, 3$, the synergy benefits of a parameter configuration are given by the difference between the information when averaged (small dot) and pooled (large dot). *Bottom* Parameter values exist where no signal combinations hold synergy (left) and synergy is equivalent across signal combinations (right).

Constrained intelligent group formation

For the groups with $k_g < 4$ (incomplete signals), there is a significantly higher variance in both information and resources compared to $k_g = 4$. This is attributed to differences in synergy between groups with different combinations of signals of order k. This illustrates a general feature of the maximal synergy principle; that signal combinations with higher conditional dependence on the environment will have higher synergy and experience higher growth rates than other combinations. Figure 4.4 demonstrates the synergy effects across different combinations of signals. For each group of size k, the left, smaller dot indicates the amount of information each signal has averaged over the signals present. The right, larger dot indicates the total information the combination of signals has when pooled. The difference between the two dots gives the amount of synergy. We see, for example, that even though signals 0 and 3 have less information than signal 2, both signals have higher synergy effects when pooled with 1 individually, as indicated by their crossover with the 1, 2 line. For a group aggregator, not only does this mean that signal choice is nontrivial, but also that individual information is not generally a good indicator of synergy benefits that can be realized when pooled.

As demonstrated by the bottom plots in Figure 4C, through a suitable selection of $\boldsymbol{p}$, we can also design special environments such as where either no synergy is present, or where there are uniform benefits of synergy across combinations of signals. The procedure for constructing environments with specific synergy profiles will be developed in future work.

These results point to the challenges of leveraging the full complementarity of available signals in practice, towards satisfying the maximum synergy principle in organizations. For example, novel signal identification may result in disruptions of existing organizational structures, which while ultimately optimal may not be realizable without some sacrifice of short-term efficiency or increased costs. Maximizing long-term synergy and growth entails a tradeoff, since over shorter horizons exploitation of existing knowledge may be preferred both

individually [157] and as organizations adapt structurally to the specialties of its members [167], which may vary depending on the complexity of the environment [241]. I have also shown that organizations that match the complexity of their environment through appropriate personnel specialization and integration (minimizing the D_{KL} with the environment) experience the fastest growth, in agreement with the analysis of empirical data [153].

4.4 Discussion

In this chapter, I developed a novel mechanism of cooperation among heterogeneous agents that use shared information to grow resources in stochastic but knowable environments. I derived the benefits of cooperation in terms of synergy gained by pooling information across agents' unique signals and its consequences for the growth rate of collectives. This motivates the principle of maximum synergy, whereby a group's aggregate growth is highest when it maximizes the synergy of its members relative to a statistical environment. I proposed this principle as a complementary avenue to cooperation resulting from the reduction of volatility through resource pooling in multiplicative growth models. I then showed that a group with no *a priori* knowledge of its potential synergy can learn it through Bayesian inference. I illustrated these principles using a model of a high-dimensional probabilistic logic gate and showed that, on average, group synergy scales superlinearly with the number of unique signals in the group. I also illustrated the challenge faced by groups incurring size-related costs to pick not just unique signals but also admit new group members as additional signals that maximize their potential collective synergy.

These results formalize several insights into the causes and benefits of cooperation. First, the formal properties of information allow us to consider how the limits to human effort and ability motivate group formation. Specialization through learning or adaptation is costly in terms of time and resources, motivating a division of labor to fully learn and maximize productivity across disparate but synergistic agents [126]. This motivates the formation of

heterogeneous cooperation networks [194, 234], where agents seek new connections that complement their particular signals and that vary conditionally on local environments. In this sense, collectively navigating complex fitness landscapes is naturally achieved by satisfying the maximum synergy principle. However, maximizing synergy can be a challenging and costly task for groups because it requires time, effort, and social rearrangement to learn the complementarities among a set of signals.

Second, these results motivate analyses of how information and resource pooling strategies affect different levels of selection within an organizational hierarchy. Effective resource pooling relies on uncorrelated fluctuations across participants, which is not possible when agents are making coordinated decisions across signals. I therefore expect information and resource pooling strategies to create tradeoffs in group formation, and apply to different environmental features and levels of selection.

Groups lacking informational complementarities (because they are homogeneous) operating in very variable environments should pool resources to minimize volatility. This may apply to people in insurance pools, or independent economic sectors within a common population, such as a city or nation. Conversely, groups in complex environments made up of agents with complementary knowledge, such as within a firm or innovation ecology, should engage in information pooling and skills specialization to maximize their collective production potential whenever the variability of the environment and costs of cooperation are sufficiently low.

Parsing out these modes of cooperation becomes more important when considering how groups respond to changing environmental or social conditions. As new environmental conditions emerge, such as new industries or technologies, the distribution of synergy across different group configurations will also change, selecting for different group compositions and skill combinations. This has the interesting implication that new knowledge (science, technology, institutional change) should be disruptive of established social and economic

structures explicitly because it enables new synergies and faster growth. This also has implications for natural ecosystems [242] where changing environmental conditions, such as via climate change, and adaptation may alter the relative fitness of their components and thereby their overall structure.

Third, the framework developed here describes a general approach to interaction dynamics in many fields. The conditional probabilities $P(e|\boldsymbol{s})$ capture the general structure of information between populations' signals and actions, and their environment. Through synergy maximization, that information becomes encoded in how groups form and are structured, and which sets of coordinated behaviors produce beneficial or detrimental outcomes across agents. By averaging over environments, we can produce a set of rules for (average) rewards associated with agents' perceptions and actions. This shows how general conditional probabilities of choices and behaviors in given environments may underlie particular "games" and other phenomenological agent interaction rules [8].

In this sense, several interesting themes in the collective dynamics of iterated games may be relatable to conditional probabilities and growth rates set by information. Two aspects of this general problem that I did not discuss here are the distribution of payoffs from collaborative action back to individual agents, and the (short-term) advantages of defection. The emergence of trust [10, 115, 150, 182] among agents necessary for realizing long-term higher growth rates is likely costly and may benefit from an aggregator that can reduce the associated risk. This catalyst of long-term synergy can also be applied to models of interaction among risky innovators [51], where coordinators can actively influence selection by managing inter-firm links and information access. In environments with conditionally dependent signals, agents may also learn to predict other agents' behavior leading to the emergence of local trust clusters [212] without the presence of an aggregator.

Thus, although the principle of maximum synergy is general, there are multiple obstacles to realizing it in practice. Pathways to explore latent synergies must overcome short-term

costs of learning and discovery, coordination, social inclusion, and exclusion, and promote the long-term bonds necessary to derive collective benefits, which once created must also be distributed fairly. When the balance of these benefits and costs is positive and can scale up, synergy becomes naturally expressed in higher-order interactions as is observed in generalized reciprocal cooperation and the emergence of complex cultures as interdependent knowledge and behavior among many agents [201].

In summary, the formal properties of information, made explicit over group structures and time, provide the theoretical basis for a broad class of agent interaction models found throughout the social and ecological sciences. This includes the formation of complex societies made up of diverse cooperating agents in situations where large-scale synergy becomes possible and can be maximized. In Chapter 6, I will discuss extensions of this theory that consider the interaction between information and resource sharing in populations across scales of organization. In the following chapter, I shift focus from cooperative to competitive dynamics.

CHAPTER 5

BELIEFS DYNAMICS IN PAIRWISE COMPETITIVE INTERACTIONS

The interaction modes considered in the previous section were strictly egalitarian, in that agents shared information, decisions, or resources towards a shared end. Some social structures, such as the hierarchical structure of a job, or the directly competitive prediction market, rely on interaction schemes that are less directly cooperative or even adversarial. In such schemes, success comes to agents who can model others' beliefs and predict their behaviors. Successful employees use signals embedded in contracts to predict the desires of their employers, prey predict predators' movements, and gamblers use the relationships between historical and current events to predict the outcomes of future events.

Predicting uncertain behavior in others is generally hard. As a result, literatures such as game and portfolio theory have emerged to help find optimal actions from a set of rules when a prediction is difficult, such as to take actions that bring competitive advantages. Theories of the mind have also been developed [202, 219], as I will touch on in this chapter, to explore more directly how agents arrive at particular decisions in complex social environments. Regardless of the setup, competitors must have some model of the preferences of agents with whom they compete in order to take a complementary action. In this sense, the dynamics of beliefs and perception become a core consideration for competitive agent adaptation. In the following section, I introduce the broad quantitative literature on studying belief formation. I then motivate why a statistical approach based in Bayesian inference is suitable for deriving preference dynamics from first principles of agent interaction. The material that follows explores the problem of belief formation using methods from nonequilibrium statistical dynamics [1].

1. This section is sourced from Kemp, J., Hongler, M. O., & Gallay, O. (2024). Stochastic pairwise preference convergence in Bayesian agents. *Physical Review E, 109(5)*, 054106.

5.1 Introduction

Belief formation is essential for studying behavior in the social and cognitive sciences. In noisy environments, empirical beliefs are formed through observation [219] and probabilistically predict future states to optimize energy costs [91]. Belief dynamics are critical for modeling how agents interact strategically in varying socio-economic contexts (through games), navigate uncertainty, and make decisions under imperfect information. However, there remain questions about how beliefs evolve in complex social environments such as networks [64] and markets [100], where the fluctuating beliefs (or perception of others') of asset values subject markets to intense volatility [151] and divergent valuations [80].

Several learning models have emerged to explain the formation of beliefs in stochastic multi-agent games [222], including frequentist and regression approaches [76, 92]. Reinforcement learning (RL) models are widely used and have intuitive descriptions [41, 202], but they do not produce closed-form solutions to dynamics of agent preferences [232], hampering the search for generalizable results. These models are generally outperformed by learning frameworks based on Bayesian inference (BI) [3, 57, 98], where agents process information to inform history-dependent, optimally predictive, and (in some cases) analytically tractable models of their environment. BI has thus become foundational in human cognition [17, 105, 147, 156], and in studying adaptive agent behavior in models of wealth and inequality [28, 141], social dynamics [187], and coordinated action [142, 248]. Additionally, Bayesian reversal learning has emerged as a more efficient alternative to reinforcement learning (RL) in more realistic, non-stationary environments [128] where discerning signal dynamics from noise is difficult [66, 73].

Solutions to closed-form belief dynamics in stationary environments have contributed to a growing literature [3, 141, 172]. However, they are not suitable for studying convergence in interacting models where signals are dynamic [127, 144]. Studying pairwise dynamics in Gaussian models, for which analytical descriptions of distribution parameters exist [178],

closes this theoretical gap while opening the door towards characterizing emergent population preference dynamics [63, 80]. We can accomplish this using established methods in non-equilibrium statistical physics, where the relationship between Bayesian inference and Ornstein-Uhlenbeck processes as noisy, mean-reverting processes with memory is well explored [185, 206, 223]. In the case of sequential Bayesian estimation, this analysis can used to study how convergence time relates to behavioral properties.

In this chapter, I propose a model for the statistical dynamics of two agents' preferences under Bayesian adaptation to another's behaviors. By treating behaviors as a Gaussian-distributed quantity, we can study the dynamics of preferences through the coupled Markov dynamics of its first-order moments. I first show that in the absence of noise, the asymptotic preferences of the agents converge to one another both to a relative value and on a timescale set by the relative strength of their priors. Later, I introduce noise and show how the dynamics resemble an Ornstein-Uhlenbeck process with time-rescaling noise. Using the Fokker-Planck equation (FPE), I then show that the preferences converge to a stationary distribution with a width set by the uncertainty in their behavior, and with dynamics governed by a relaxation time, $t^\star$. I conclude by discussing how convergence can be broken by introducing unpredictable behavioral shocks, and the model's implication for studying belief formation in a host of game-theoreticl and principal-agent problems.

5.2 Bayesian Preference Dynamics

Consider agents A and B, who at time-step i exhibit a statistically distributed, real-valued behavior $x_i \in X$ and $y_i \in Y$. I denote the normalized distribution of their decisions $P_i(X|\theta_A)$ and $P_i(Y|\theta_B)$, parameterized by behavioral parameters θ_A, θ_B. Consider that the agents can learn each other's behavior and are motivated to align their decisions (e.g., $[x_i - y_i]^2$ is minimized), but cannot directly coordinate their actions before observation. While coordination can be accomplished by conditioning behavior on some shared signal [142], this would not

change the general dynamics and is excluded for brevity.

Each agent infers the other agent's preferences by observing their cumulative noisy behavior and adjusting their preferences to match. The term preferences implies the first moment of the distribution of behaviors that spans the agent's set of choices. This particular setup is motivated by open questions in principal-agent problems, where agents must coordinate their behavior through adaptation [214].

History-dependent learning is accomplished optimally through BI [141]. As such, the distribution of agent A's behaviors at $i = 0$ forms a prior for their guess of $B's$, $P(X = x) \equiv P_0(\tilde{Y} = x)$, for approximated behavior $\tilde{Y}$ (and $\tilde{X}$ for B). The distribution of decisions at later interactions is given by a posterior $P_i(\tilde{Y}|\{y_i\})$, where the decision is conditioned on the history of B's behavior [2]. After n steps, agent A's posterior is given by (and B by analogy)

$$P_n(\tilde{y}|\{y_i\}, \theta_A) = \left[\prod_{i=1}^{n} \frac{P(y_i|\tilde{y}_i)}{P(y_i)}\right] P(x|\theta_A), \tag{5.1}$$

In sequential Bayesian inference, an agent's behavior at step n follows a Markov process and is sampled from P_{n-1}. This process is illustrated in Figure 1, where $\mathcal{L}$ denotes the likelihood given the evidence.

In this work, I assume the behaviors are instantaneously described by Gaussian distributions with gamma-distributed priors, $x \sim \mathcal{N}(\mu_x, \sigma_x|\boldsymbol{\theta}_x)$ and $y \sim \mathcal{N}(\mu_y, \sigma_y|\boldsymbol{\theta}_y)$, where $\boldsymbol{\theta}$ is the gamma prior vector. The means, μ_x, μ_y, describe the agents' preferences, whereas the fluctuation in true behavior is given by the Gaussian standard deviations σ_x, σ_y.

Bayesian inference on this choice of distribution results in preference dynamics that are linear [178]. Therefore, I first study the dynamics of the preference averages, then later consider how noisy behavior couples into the preference variances. The following analysis gives a first-order approximation of the complete behavior (*vis a vis* the preferences) un-

2. Through observation, $\tilde{Y}$ is conditioned on $\tilde{X}$, and agent A's past behavior, filtered through B, influences their future behavior. B is similarly conditioned by A's.

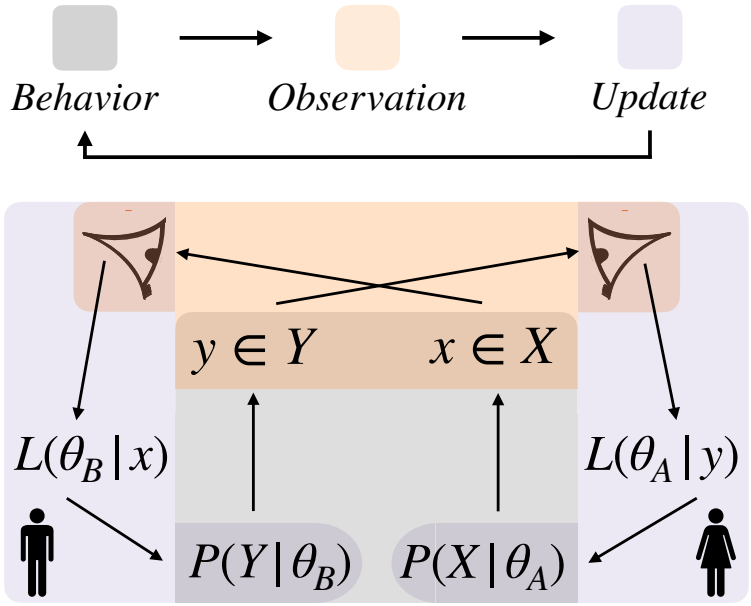

Figure 5.1: Diagram of the interaction model. Agents A and B sample behaviors $x \in X$ and $y \in Y$ from respective distributions. Agent A updates their prior θ_A with the evidence $\mathcal{L}(y|\theta_a)$ from B's behavior, and vice versa.

der Bayesian inference, whereas dynamics of higher order naturally come from higher-order moments and their couplings. This study assumes $\sigma_x = \sigma_y$ and leave the dynamics of the standard deviations under Bayesian inference for future work.

5.2.1 Deterministic Dynamics

First, I will study the dynamics of the preference parameter in the absence of noise. The rule for updating the mean parameter of a Gaussian-Gamma model under Bayesian inference is described recursively after n steps as (APP A.4.1) [178]

$$\mu_x^n = \frac{\mu_x^{n-1}\left(\frac{n-1}{\omega} + \alpha\right) + \mu_y^{n-1}}{\frac{n}{\omega} + \alpha}, \quad \mu_x^1 = \frac{\alpha x_0 + y_0}{1 + \alpha},$$

where $\omega = n/t$ is the interaction rate. In the continuous limit $\omega \to \infty$, μ_x and μ_y become coupled by the linear differential equations

$$\frac{\partial \mu_x}{\partial t} = \frac{\mu_y(t) - \mu_x(t)}{t + \alpha}, \quad \frac{\partial \mu_y}{\partial t} = \frac{\mu_x(t) - \mu_y(t)}{t + \beta}, \tag{5.2}$$

where x_0, y_0 are the initial preferences and α, β are the hyperprior magnitudes with units t. Denoted the learning times, these parameters measure how resilient the preferences are to new evidence. These equations say that the dynamics of the preference parameters μ_x, μ_y decrease as the quantities converge in time. I demonstrate this by constructing the ODE for the difference measure $\Delta(t) = \mu_x(t) - \mu_y(t)$ (correspondingly $\Sigma(t) = \mu_x(t) + \mu_y(t)$), with solution (APP A.4.1)

$$\Delta(t) = \frac{\Delta_0 \alpha \beta}{R(t)}, \tag{5.3}$$

where $\Delta_0 = x_0 - y_0$, and $R(t) = (\alpha + t)(\beta + t)$ is the time rescaling coefficient. This shows intuitively that the agents' preferences converge with power law -2 in time that increases symmetrically as $\alpha, \beta \to \infty$, and agent learning times increase.

With intuition for the coupled system established, we can now study the dynamics of the full system. There exist two solutions to Eq. 5.2 given by the equality of the learning times. First, when $\alpha = \beta$, the dynamics have the asymptotically symmetric solution $f(x_0, y_0, t) = \mu_x(t)$ and $f(y_0, x_0, t) = \mu_y(t)$, where f is defined as

$$f(x_0, y_0, t) = \frac{2\alpha^2 x_0 + (2\alpha t + t^2)(x_0 + y_0)}{2(\alpha + t)^2}. \tag{5.4}$$

It follows that $\lim_{t \to \infty} f = (x_0 + y_0)/2$, and both agents' preferences converge to the average of their initial preferences asymptotically, at times $t \gg 2\alpha$.

In the case $\alpha \neq \beta$, the solution for μ_x is given by

$$\mu_x(t) = \frac{\alpha x_0}{\alpha + t} + \frac{\alpha\beta(x_0 - y_0)}{(\alpha - \beta)^2} K(t) + \frac{t(\alpha x_0 - \beta y_0)}{(\alpha - \beta)(\alpha + t)} \tag{5.5}$$

where $K(t) = \ln\left[\frac{\alpha\beta+\beta t}{\alpha\beta+\alpha t}\right]$ is a dynamical value with $\lim_{t\to\infty} K(t) = \ln[\beta/a]$. As we would expect, $\mu_x(0) = x_0$ and at long times, $\mu_x(t) \to s$, where s is the weighted average between the initial values,

$$s \equiv \frac{x_0\big[\alpha/\beta + \ln[\beta/\alpha] - 1\big] + y_0\big[\beta/\alpha - \ln[\beta/\alpha] - 1\big]}{(\alpha-\beta)^2/\alpha\beta}.$$

The solution for $\mu_y(t)$ is given in the appendix, with $\mu_y(t) \to s$ asymptotically. These results are demonstrated at the top of Fig. 5.2 for various learning times, with $y_0 = 5$ and $x_0 = 1$. In matrix form, these dynamics are given by $\mathbb{M}[x_0, y_0] \equiv [x(t), y(t)]$, where the drift matrix is

$$\mathbb{M} = \frac{\alpha\beta}{\alpha-\beta}\begin{pmatrix} M_2(t) - M_1(0) & M_2(0) - M_2(t) \\ M_1(t) - M_1(0) & M_2(0) - M_1(t), \end{pmatrix}$$

$$M_1(t) = K(t) - \tfrac{1}{(t+\alpha)}, \;\; M_2(t) = K(t) - \tfrac{1}{(t+\beta)}.$$

This invertible matrix has a nonzero determinant $\det[\mathbb{M}(t)] = \alpha\beta/R(t)$. As we will see, this gives the constant of motion for constructing exact solutions for the dynamics of the system with noise [49].

Asymptotic preference behavior

Conveniently, the asymptotic preference value can be expressed independently of the initial condition, allowing us to compute the relative shift in preferences as a function of learning times. Consider the initial parameter difference Δ_0, and the difference in asymptotic value from the initial parameter $\delta_x = x_0 - s$. The fractional similarity of X is given by $f_x = 1 - \delta_x/\Delta_0$. This expresses how close X has remained to x_0 relative to y_0, and is useful for measuring the change in preferences of an agent represented by X (and Y by analogy). It is

given by

$$f_x = 1 - \alpha\beta\frac{\beta/\alpha - \ln(\beta/\alpha) - 1}{(\alpha-\beta)^2}, \;\; f_y = 1 - f_x. \tag{5.6}$$

These fractional limits are demonstrated in the bottom of Fig. 5.2 over various learning times.

So far, I have explored the dynamics of this model without noise, and have shown that both preference parameters converge to a value set by the relative magnitude of the learning times. I have shown that the deterministic dynamics are isomorphic and that we glean useful information about the relative change in preference between the agents without knowledge of the initial conditions. These results establish intuition for how, on average, agent characteristics determine the convergence process. In the following section, I will introduce noise to the inference process, and demonstrate a procedure for constructing exact solutions using the linear and isomorphic properties of the dynamics. While this procedure results in lengthy analytical solutions that are not explored, I will demonstrate some key insights from the coupled dynamics, $\Delta(t), \Sigma(t)$.

5.2.2 Full Dynamics under Noisy Sampling

I introduce noise by rewriting Eq. 5.2 as the stochastic differential equations on quantities X_t, Y_t,

$$dX_t = \frac{\sigma_y}{t+\alpha}dW_{2,t} - \frac{\Delta_t}{t+\alpha}dt, \;\; dY_t = \frac{\sigma_x}{t+\beta}dW_{1,t} + \frac{\Delta_t}{t+\beta}dt,$$

with boundary conditions $X_0 = x_0, Y_0 = y_0$. I have introduced white Gaussian noise (WGN) processes, $dW_{1,t}$, $dW_{2,t}$ with magnitudes σ_x, σ_y that describe i.i.d fluctuations in agent behavior. Recalling previously that the asymptotic preferences depend on the initial conditions,

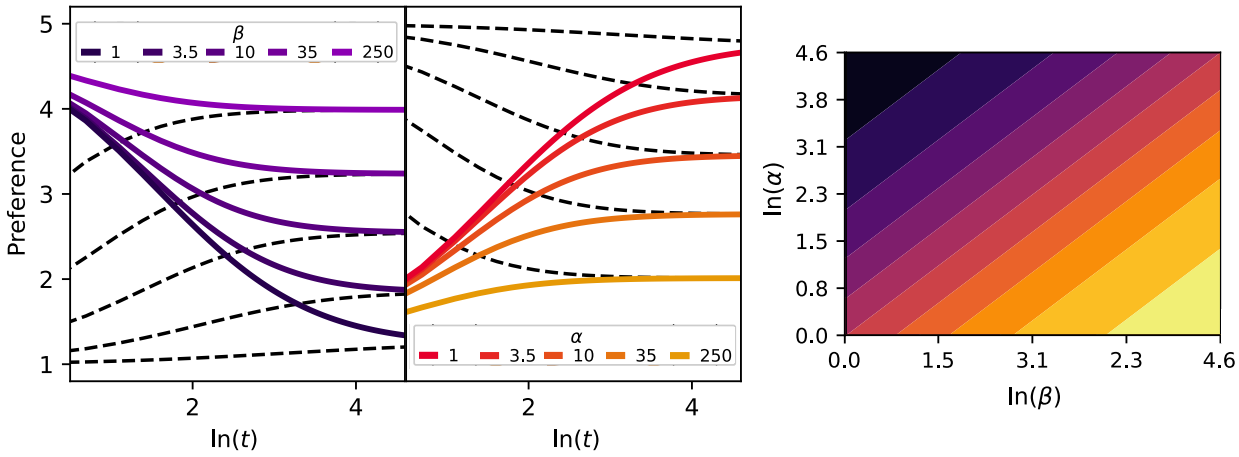

Figure 5.2: Behavior of the noiseless model in Eq. 5.2 with $x_0 = 1, y_0 = 5$. *Top*: Convergence values computed from Eq. 5.5 for variable β (left) and variable α (right) for constant reference agent ($\alpha, \beta = 5$), represented by dashed lines. *Bottom:* Asymptotic fractional drift computed from 5.6 on a logarithmic parameter scale.

Note that while the dynamics of the SDEs are Markovian, they cannot be ergodic. The dynamics of both preferences behave like Ornstein-Uhlenbeck (OU) processes, as the magnitude of the attractive drifts increases with the magnitude of the difference. However, this OU process is time inhomogeneous, as the magnitude of all dynamics decay with a power law in time. We interpret these dynamics in terms of the underlying Bayesian inference process. The rate of parameter convergence slows as the agents converge in parameter value, and the effect of each interaction decreases in time as the agent weighs cumulatively larger sums of evidence. At long times, when preferences have nearly converged and have accumulated lengthy histories, small fluctuations dominate the dynamics.

To explore the statistics of the two-dimensional process, I define the bivariate transition probability distribution (TPD) as $P(x, y, t|x_0, y_0)$. The evolution for this distribution is given by the FPE, $\partial_t P(x, y, t|x_0, y_0) = \mathcal{F}\big[P\big]$, where $\mathcal{F}[\cdot]$ is defined

$$\begin{aligned}\mathcal{F}[\cdot] =& \partial_x\big(\tfrac{y-x}{t+\alpha}[\cdot]\big) + \partial_y\big(\tfrac{x-y}{t+\beta}[\cdot]\big) \\ &+ \tfrac{\sigma_y^2}{2(t+\alpha)^2}\partial_{xx}[\cdot] + \tfrac{\sigma_x^2}{2(t+\beta)^2}\partial_{yy}[\cdot].\end{aligned} \tag{5.7}$$

One can marginalize the distribution for x, $P_M(x,t|x_0) = \int_{\mathbb{R}} P(x,y,t|x_0,y_0)dy$, and by analogy, y. To solve these equations exactly, I transform the set of equations into the frame of constant motion, defined by $\mathbb{M}(t)$, in which the process is purely diffusive and described by a Gaussian. In this frame, solutions for the dynamics of $P_M(x,t)$ and $P_M(y,t)$ are exactly solvable [49]. However, this procedure does not lead to concise results and is detailed only in the appendix.

As in the deterministic case, we glean tractable insights into the dynamics by solving the FPE for the coupled system, $X_t \to \Delta_t = X_t - Y_t$, $Y_t \to \Sigma_t = X_t + Y_t$. In the following section, I will use an exact solution of the FPE to show how the mean and variance of the TPD of Δ_t converges to zero, encoding the system's entropy into Σ_t. I will conclude this work by approximating an upper bound for the asymptotically stationary variance of $\Sigma(t)$.

5.2.3 Solutions of the FPE for Coupled Dynamics

In terms of the original model parameters, the new SDEs are

$$\begin{aligned} d\Delta_t &= \frac{\sqrt{\sigma_y^2(t+\beta)+\sigma_x^2(t+\alpha)}}{R(t)} dW'_{1,t} - \frac{2t+\alpha+\beta}{R(t)}\Delta_t dt \\ d\Sigma_t &= \frac{\sqrt{\sigma_y^2(t+\beta)+\sigma_x^2(t+\alpha)}}{R(t)} dW'_{2,t} + \frac{\alpha-\beta}{R(t)}\Delta_t dt, \end{aligned} \tag{5.8}$$

where the dW' terms are now correlated White Gaussian noise processes. Again, we see that the difference equation behaves like an OU process, where drift is set by the difference in preferences, with time-rescaling noise. In this sense, Σ_t does not couple into the dynamics of Δ_t, permitting us to solve for the statistics of Δ_t first, then Σ_t.

The Difference Equation

In these coordinates, the statistics of Δ_t are fully described by the TPD $P_\Delta(z,t|z_0,0) = \text{Prob}\{z \leq \Delta_t \leq (z+dz)|z_0\}$, which solves the FPE

$$\partial_t P_\Delta = \partial_z\big[\tfrac{2t+\alpha+\beta}{R(t)} z P_\Delta\big] + D(t)\partial_{zz} P_\Delta, \tag{5.9}$$

where the diffusivity $D(t) = \frac{\sigma_y^2(t+\beta)^2+\sigma_x^2(t+\alpha)}{2[R(t)]^2}$. To solve this partial differential equation, we seek the reference frame where the process becomes purely diffusive. Consider the change of variables $z \mapsto z' \equiv z\frac{R(t)}{\alpha\beta}$, and $t \mapsto \tau \equiv t$. The differential operators transform as $\partial_z \mapsto \frac{R(t)}{\alpha\beta}\partial_{z'}, \partial_t \mapsto \frac{2t+\alpha+\beta}{R(t)} z'\partial_{z'} + \partial_t$, where I used the equivalence $t = \tau \to \partial_t = \partial_\tau$, yielding $\partial_t P_\Delta = \frac{2t+\alpha+\beta}{R(t)} P_\Delta + D(t)\frac{R(t)^2}{\alpha^2\beta^2}\partial_{z'z'} P_\Delta$ (APP A.64). Introducing the rescaling $P_\Delta \equiv R(t)Q_\Delta$, diffusion absorbs the drift term and reduces the dynamics to time inhomogeneous diffusion $\partial_t Q_\Delta = D'(t)\partial_{z'z'} Q_\Delta$ where $D'(t) = \frac{\sigma_x^2(t+\alpha)^2+\sigma_y^2(t+\beta)^2}{2\alpha^2\beta^2}$. To solve this equation, I introduce the time rescaling, $t \to s(t) = \int_0^t D(\xi)d\xi$, giving

$$s(t) = \frac{\sigma_x^2(t+\alpha)^3 + \sigma_y^2(t+\beta)^3}{3\alpha^2\beta^2} - s_0,$$

where $s_0 = \frac{\sigma_x^2\alpha}{\beta^2} + \frac{\sigma_y^2\beta}{\alpha^2}$. Eq. (5.9) has a Gaussian solution $Q_\Delta \sim \mathcal{N}\big(z_0', \sqrt{s(t)}\big)$. This Gaussian transforms back to the moving frame, $z' \to z = z'\alpha\beta/R(t)$ to get the full solution for P_Δ

$$P_\Delta(z,t|z_0,0) = \frac{R(t)}{\sqrt{2\pi\sigma_\Delta^2(t)}} \exp\Big[-\frac{\big(z - z_0\frac{\alpha\beta}{R(t)}\big)^2}{2\sigma_\Delta^2(t)}\Big],$$

where I have introduced the difference variance as

$$\sigma_\Delta^2(t) \equiv \frac{\sigma_x^2(t+\alpha)^3 + \sigma_y^2(t+\beta)^3}{3R^2(t)}. \tag{5.10}$$

This probability density has a few key features. First, $\lim_{t\to\infty}\langle\Delta_t\rangle = 0$ and $\lim_{t\to\infty}\sigma^2_\Delta(t) = 0$ so the distribution asymptotically converges to a delta function at $\Delta = 0$. By construction, $\sigma_\Delta(0) = 0$ and $\sigma^2_\Delta(t) \geq 0$, meaning the variance evolves non-monotonously, and maximizes at a relaxation time $t^\star$.

While cumbersome, these results are intuitive in terms of the underlying inference process. Agents' preferences converge to one another almost certainly in a time that increases with the magnitude of fluctuations but decreases in the strength of the agents' learning times. Therefore, the strength of attraction is asymptotically stronger than noise fluctuations. These convergence dynamics are demonstrated by Monte Carlo (MC) simulations in Figure 5.3 with $N = 1000$, $\alpha = \beta = 15$, where we see that $\Delta \to 0$ on a logarithmic scale in agreement with theory. Define the observables $m(t) = \Sigma(t)/2$, $\Delta_m(t) = m(t) - m(0)$, where Δ_m measures how the mean of the dynamics change over time, and $\Delta_{y,m} = m(t) - y(t)$ demonstrates how y differs from the mean. We see that $\Delta_{y,m}$ ($\Delta_{x,m}$) converges to zero with asymptotically vanishing noise, indicating that the preferences are converging to the mean almost certainly, while the entropy is increasingly expressed through the convergence value.

However, convergence in realistic agents are subject to social, environmental, and biological constraints to their interaction time (such as life expectancies). So, while convergence is asymptotically guaranteed, it is not guaranteed for particularly noisy or stubborn agents that have interaction time horizons shorter than the convergence timescale.

In the following section, I will use these insights to solve for the TPD of the Σ_t process.

The Sum Equation

The cumbersome dynamics of Δ_t lead to an even more complex analytical description for Σ_t. However, we know that $\lim_{t\to\infty}\langle\Delta\rangle(t) = 0$ and $\lim_{t\to\infty}\sigma_\Delta(t) = 0$. It follows that the initially bi-variate diffusion process asymptotically collapses into a uni-variate pure diffusion centered at $\mu_{\Sigma,f}$. Hence, for $t \to \infty$, the transition probability is approximated as [139]

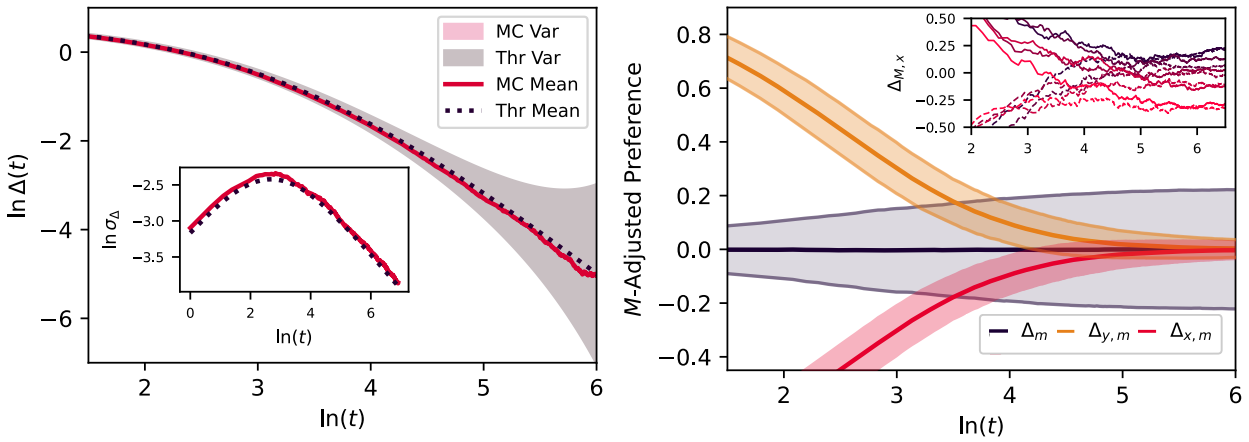

Figure 5.3: Coupled dynamics of stochastic agent preferences with $x_0 = 0, y_0 = 2$, $\alpha = \beta = 25$ and 95% CI shaded regions. *Top:* MC Dynamics of Δ_t match theory. *Inset:* The variance initially increases, reaches a maximum at $t^\star$, and then decreases. *Bottom:* The mean-adjusted dynamics ($\Delta_m = \Sigma/2$) is constant for this choice of parameters, with asymptotically constant noise. The difference in agent parameters from the mean, $\Delta_m - y$, $\Delta_m - x$ converge to 0 with time vanishing noise. *Inset*: Selected trajectories demonstrating different asymptotic values.

$$d(\Sigma_t - \mu_{\Sigma,f}) = d\Sigma_t \approx \frac{\sqrt{\sigma_y^2(t+\beta)+\sigma_x^2(t+\alpha)}}{(t+\beta)(t+\alpha)} dW'_{2,t},$$

where the constant $\mu_{\Sigma,f}$ is given in Eq.(A.55). The corresponding TPD $P_\Sigma(z,t|\mu_{\Sigma,f})dz = \text{Prob}\{z \leq \Sigma_t \leq (z+dz)|\sigma_0\}$ solves the FPE:

$$\partial_t P_\Sigma = \frac{\sigma_x^2(t+\beta)^2+\sigma_y^2(t+\alpha)^2}{2R(t)} \partial_{zz} P_\Sigma.$$

By inspection, P_Σ is a Gaussian law with mean $\mu_{\Sigma,f}$ and by an *ad-hoc* time re-scaling, the time-independent variance is computed as [139]

$$\langle \Sigma_t^2 \rangle = \sigma_x^2 \left[\frac{1}{\alpha} - \frac{1}{\alpha+t}\right] + \sigma_y^2 \left[\frac{1}{\beta} - \frac{1}{\beta+t}\right], \tag{5.11}$$

We now see that the second moment (and hence the variance) converges in the limit $t \to$

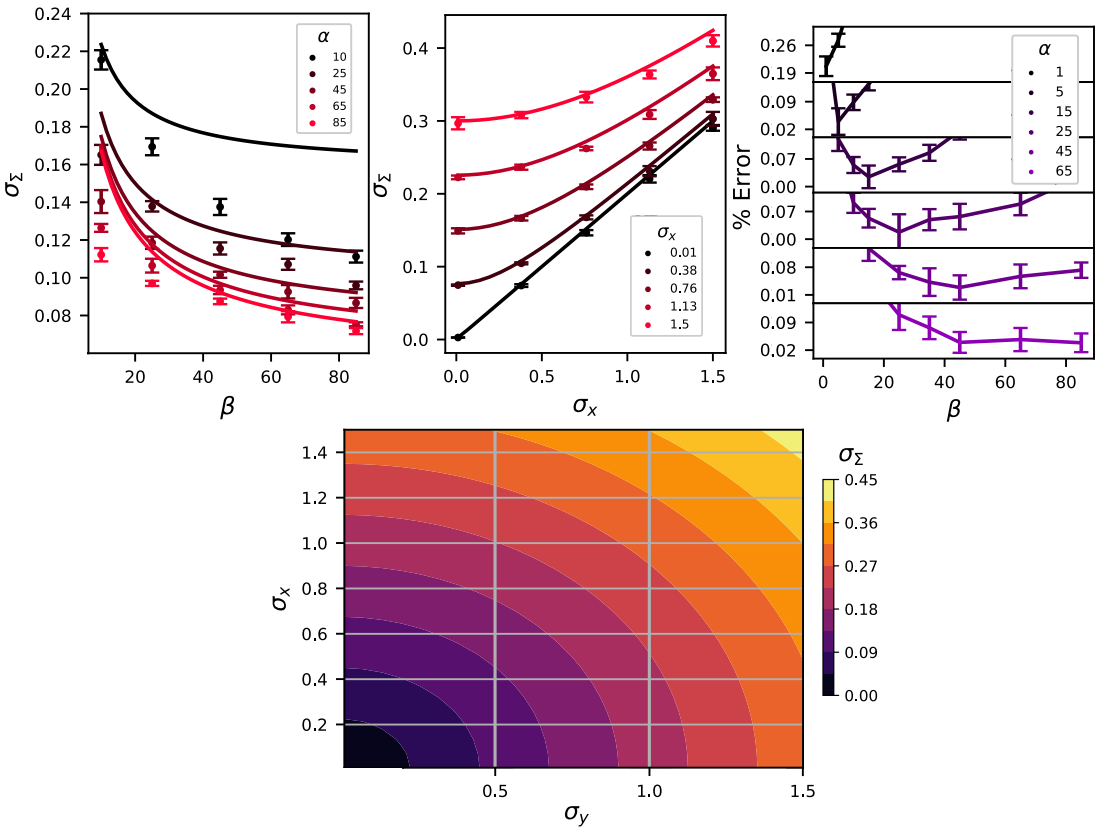

Figure 5.4: Asymptotic variance upper bound $\langle \Sigma_t^2 \rangle$ under various parameters *Top:* $\langle \Sigma_t^2 \rangle$ increases with agent noise with agreement between MC simulations and theory. *Bottom:* Variance decreases with agent learning time. *Left:* Variance upper bound $\langle \Sigma_t^2 \rangle$, diverges from empirical results as β and α diverge. *Right:* Deviation between the upper bound and MC experiments as a fraction of theoretical prediction.

∞ to a stationary value $\langle \Sigma_t^2 \rangle = \frac{\sigma_x^2}{\alpha} + \frac{\sigma_y^2}{\beta}$. Similarly to the case of σ_Δ^2, the variance of this distribution increases with the fluctuations in behavior and decreases with hyperprior strength. Note, though, that this is an upper-bound estimation for variance, as [139]

$$\sigma_\Sigma^2(t) = \langle (\Sigma_t - \mu_{\Sigma,f})^2 \rangle = \langle \Sigma_t^2 \rangle - \mu_{\Sigma,f}^2 \leq \langle \Sigma_t^2 \rangle.$$

Furthermore, from the convergence of $\Delta_t \to 0$, we know preferences converge to $X_t = Y_t$, and the variances converge to $\langle X_t^2 \rangle = \langle Y_t^2 \rangle = \frac{\langle \Sigma_t^2 \rangle}{2}$ as $t \to \infty$.

The asymptotic preference variance demonstrated in Eq. 5.11 shows that more noisy agents who learn quickly will converge to more entropic states. This is because the fluctuations of one agent are recorded and immediately reciprocated by the partner agent (on average), biasing future preferences towards early fluctuations. When noise is strong or learning fast, agents weigh fluctuations more heavily than weak noise or slow learning. In both cases, these characteristics more strongly ossify early, noisy shifts in preferences before mean behavior is fully resolved, coupling more noise into the asymptotic behavior of the system. MC simulations demonstrate the relationships between these quantities and the variance in Fig. 5.4. We also see through simulations that the percent error in the upper bound estimate and the true variance are closest when $\beta = \alpha$, and increases as the learning times diverge. However, this difference decreases as both learning times approach large values.

5.3 Discussion

In this chapter, I studied a simple model for pairwise belief formation in Bayesian agents who adapt to each other's behaviors. I showed that preferences converge on a timescale and to a value given by the agents' relative learning times. Using the Fokker-Planck equation I then explored the convergence characteristics of the Gaussian PDF for preferences in the combined frame. I showed that while agents' preferences invariably converge to one another,

the relative value is noisy, is characterized by a relaxation time $t^\star$, and is bounded above by a sum of the standard deviations of agent behaviors weighted by the learning times.

There remain several challenges concerning the full characterization of this system. First, deriving the full dynamics of Σ_t would be useful for attaining better bounds on the asymptotic coupled behavior. Second, solving the nonlinear dynamics for the covariance matrix of the inference process gives the full dynamics of the interaction, although this would likely require numerical treatment. Once we understand the full dynamics of this interaction, we can scale this model to include agents with multidimensional, co-varying preferences. This builds towards a Bayesian analog of a Self-Other Model [202], wherein agents coordinate decisions by approximating the other agent's behavior and serve as a microfoundation for more robust statistical mechanical models of network belief formation [63, 64]. However, doing so also requires mechanisms for polarization, such as biased assimilation [162] where agents become resistant to preferences that diverge strongly from their own. While this behavior has been explored in DeGroot models of opinion formation [65], I must still explore how these interactions compete with convergence dynamics in a Bayesian context.

Further work can extend this analysis by studying preference dynamics in agents that must balance learning each other's signals with some additional, external signals. When only one agent observes an additional, stationary signal, it is natural that the agents' preferences would converge to a value biased by the external signal. However, when one agent observes an additional, non-stationary signal, as a form of unpredictable shock, or both agents observe separate, stationary signals as a form of "reality check", their preferences are not guaranteed to converge [80]. The existence of a phase transition would depend on whether the external signal alters their preferences on a timescale comparable to the relaxation time $t^\star$, and can be applied at scale to study many-body preference dynamics.

Although, this work can already be applied to game theoretic models of dynamical persuasion [134], dynamical prisoner's dilemma [154] and other games of trust and coordination

[10, 51, 149], where the transmission of preferences through behavior determines asymptotic Nash equilibria. This formalism can also be adapted to study how preferences evolve in principal-agent models where there remain questions of how the value and variance of asymptotic preferences behave as agents adapt to post-contract disagreements. In cases where the agent serves as an information channel for the principal [214]. Models of information-driven resources dynamics [141] can be used to study how the convergence rate affects agent resources, and how the entropy of convergence values affects the quality of information transmitted to the principal. Generally, these results constitute a step towards more robust quantitative models of inter-agent and market interactions that incorporate findings from the cognition community.

CHAPTER 6

CONCLUSION AND OUTLOOK

In this dissertation, I derived general nonequilibrium statistical dynamics for the behavior of adaptive agents in uncertain environments that reward accurate prediction. I have shown that agents optimally grow resources in proportion to the information between their predictive signal and the future states of their environment. Information theory then provides a natural framework to study adaptive and interactive decision-making processes. For example, growth is maximized in the long run through Bayesian inference, as agents use their past experiences in their environment to improve decision-making given their signal. Other decisions, such as the selection of signals or with whom to pool information are described by the availability and compatibility of information across signals. These results demonstrate the deep connections between the emergence of growth, inequality, cooperation, and competition among intelligent, forward thinking agents. This work provides a highly modular starting point for building models of behavior on dynamical networks, based on principles of optimal learning, using a minimal number of assumptions about the structure of agent behaviors or preferences, and environments.

In this work, I have primarily shown what decisions are optimal and how to learn from experience optimally, given a signal or set of signals. To apply these principles to real-world dynamical systems, I must now turn to understanding how suboptimal agents navigate their environments. We must study how agents evaluate the expected payoffs of a signal and how that computation affects their signal selection. We also must understand how agents project the payoffs of learning, and how environmental characteristics such as volatility or life expectancy, may affect their decision to invest in learning. Lastly, we must develop dynamics for group formation given the principles of optimality in group behavior developed in this work. That is, understand how the availability of synergy in an environment affects agents' decisions to form collectives under costs to sharing information. Understanding these

concepts will allow us to build models of adaptive, collective dynamics that avoid heuristics or unverified assumptions, and build towards a more unified understanding of complex systems.

Achieving this goal involves applying control theory guided by principles developed in this work. Control theory provides principled mechanisms for understanding optimal adaptation in nonstationary, noisy environments with feedback from the previous decisions of agents. This allows us to understand how individuals learn under uncertainty to the value of payoffs, and how emergent agent collectives satisfy the principle of maximum synergy. In this section, I will define the key challenges to reaching this goal, and then illustrate three general research opportunities.

6.1 Adaptation in volatile environments

The mathematics of compounding dynamics show that volatility (growth rate fluctuations) reduces the rate of returns to growth and learning [140]. Empirical evidence in psychology suggests high volatility constrains agents to exploit more short-term gains rather than exploring [90], leading to higher discounting of possible future gains. It remains unclear in this context to what extent agents can arrive, over a life course, at optimal policies, consistent with their environmental opportunities. This raises several questions: In terms of adaptation, does high volatility early in life (or in a learning process), or evidence of a low life expectancy (through environmental evidence) contribute to strategies that are more exploitative and less exploratory? Do agents in these conditions adhere to the long-term optimal strategy described in prior work, or do agents rescale their optimization time? Finally, can we determine when is the theory not predictive at either the level of populations or individuals, and can we identify what behaviors require new theoretical mechanisms?

To address these questions, we require learning techniques, such as reinforcement learning (RL), that describe adaptation under constraints to exploration such as learning costs or budgeting. Through a combination of RL techniques guided by the information-driven

resource and cooperative dynamics described in prior work, we can study how agents in these constrained environments learn and develop policies dynamically. This will determine whether decision parameters evolve predictably on average by comparing learning histories between the agents described by the theory in this dissertation and RL agents, and experimentally verify theoretical schemes of resource-constrained learning and group formation. In general, this will enable a comparison between these policies and the theoretical optimal, how policy disparities contribute to inequality across agents and environments, and whether collectives of agents under these constraints realize the maximum synergy principle. This will clarify in which environmental circumstances agents fail to learn the available information in their environment, arriving at divergent information or suboptimal cooperatives.

This work would contribute to our understanding of adaptation in highly volatile environments and make several contributions to advances in theories of social choice and artificial intelligence (AI). It will improve our understanding of noisy cognition, a topic critical to studying learning, growth, and inequality in realistic settings where agents make intertemporal decisions. This work also motivates and contributes to the study of learning under non-ergodic reward schemes, an emerging topic in the AI and living systems literature [13].

6.2 Tradeoffs in Cooperative Strategies

Agents that make collective decisions over shared signals experience synergy benefits that depend on the conditional dependence of their signal given an event. I established in this text that information and resource sharing provide complementary modes of cooperation. However, I have not yet characterized the nontrivial tradeoffs between these two strategies depending on description of the environment. Particularly, signals that are more statistically independent tend to promote resource sharing, while statistically dependent signals may promote information sharing, depending on their conditional dependence. Furthermore, agents that make conditionally independent decisions, and thus experience statistically in-

dependent growth in a stochastic environment, benefit from resource sharing on timescales that grow with the environment's volatility [79]. This trade-off is reflected in the varied strategies observed in human organizations. For example, firms that prioritize innovation and market capitalization strive to integrate new signals through research and development or recruitment, in a sense maximizing their synergy. Conversely, governments pool resources to redistribute wealth or provide service, through tax solicitation. Some collectives like conglomerates engage in a mixed cooperation strategy by overseeing diverse portfolios of firms, a strategy that hedges uncertainty between firms while leveraging synergy within firms. There remain open questions about whether there is an optimal cooperation strategy tied to environmental characteristics, whether there exist discrete phase transitions between dominant cooperation strategies, and under which circumstances mixed strategies emerge when there are costs to resource and information sharing. Additionally, we need theoretical frameworks that can handle how these considerations differ across scales of cooperation.

These questions can be answered through a direct application of the methods developed in this dissertation, both standard statistical physical treatments and numerical simulations. Previous literature [79] established a suitable order parameter that measures the fraction of individual resources shared between members of a population. Analogously, we can define an information channel between agents within a group with tunable noise parameters. Scaling the noise parameter controls how well a signal is transmitted between members of a group, making uncertain the benefits of signal sharing under transmission costs. We can then compute the sharing benefits, terms of synergy, as a function of the noise parameter, and associated sharing costs, identifying the threshold cost to sharing an agent can bear. This analysis can then be repeated for the combined case of information and resource sharing to produce a two-dimensional phase diagram. This will illustrate which strategy, between information sharing (affected by transmission noise), and resource sharing (affected by signal coupling) emerges as the dominant strategy. These results can be validated through Monte-

Carlo simulations of dynamics similar to those in prior works, and compute time horizons for realizing the average growth rates.

These calculations provide a critical development towards constructing fully integrated models of agent resource growth in cooperative environments. This will not only improve the existing toolkit for studying inequality and cooperation phenomena in stochastic multiplicative growth models but will take a necessary, theoretically sound step towards modeling heterogeneous interactions and games in more complex spatial networks.

6.3 Principal-Agent and Other Competitive Interactions

The value of information in firms and markets has been well known for some time [118]. It has been studied in the context of financial markets, where decisions are often optimized for immediate returns, but not always. This market information is aggregated by averaging competing signals that each introduce insight about the value of some product. However, the dynamical effects of other aggregation methods such as pooling have not yet been studied in market contexts. As a result, there is a critical lack of theory regarding market dynamics among forward-thinking, cooperative agents optimizing returns in a noisy environment. Particularly, there are open questions about whether agents' preferences converge in the presence of external signals [80], and whether adaptive agents provide stable information channels under pressure from management [215]. Furthermore, positive feedback systems are present in realistic environments whereby agents' behaviors influence the time evolution of the statistics of their environment. The mechanisms of these interactions remain unexplored, as are the effects of these interactions on the emergence of inequality and social organization in socially driven environments.

Previous work on preference convergence [139] introduced an analytical approach to studying competing Bayesian agents, a first step toward understanding multi-agent interaction schemes. Using stochastic dynamical theory supported by Monte Carlo simulations,

we can explore how external, stationary signals, representing a "reality check" and non-stationary signals, representing generalized shocks, affect the convergence properties of the agents' preferences. Similar to the prior topic, we can explore whether statistical phase transitions exist between convergent and non-convergent preferences across populations on finite timescales. We can then fit this model to fluctuating market data and introduce market resource dynamics [28] (i.e. information advantages translate to growth, rather than total information) to study the evolution of resource inequality in market contexts. We can also apply these results to ongoing research into the dynamical principal-agent problem among agents that balance various internal and external signals [215]. This research will use similar resource and information dynamics to explore how the adaptability (via learning rates) of principals and agents alike affect the long-term performance (via resources) of public institutions.

This research quantitatively incorporates results from the cognitive science community and stochastic growth dynamics to address questions in the organizational sciences and public performance literature. Particularly, it will enable principled study of reputation dynamics in the principal-agent problem and the evolution of preferences in dynamical markets. This development will contribute to a broad literature on the properties and dynamics of markets. Specifically, it will enable us to quantitatively articulate how the dynamics of individual growth or the complementarity of skills and information across competing organizations affect the dynamics of prices and the emergence of inequality in competitive settings.

6.4 Expanding the environment

The previous three sections have deepened our understanding of agent adaptation given an environment. We accomplished this using simple, stationary, black-box sampling of events conducive to closed form expressions for learning and resource dynamics. While these treatments gave us valuable, general insights, real-world environments are more complex. Incor-

porating these complexities can advance our understanding of agent adaptation.

We have identified four key mechanisms that should be studied to enrich our understanding of adaptive agent behavior. First, environments often exhibit history dependence, where past states and decisions affect future conditions. Second, the sampled state of an environment may consist of several potentially causally related events, to which agents with different functions respond uniquely. Third, agents experience feedback through their environment from their decisions and decisions of others, shaping future decisions. Fourth, environments have internal constraints, like energy availability, and they impose constraints on agents, such as life expectancy. Although these considerations are not independent, each imposes unique constraints that break the stationarity assumption, affecting the resource optimization and learning processes. Addressing these factors is critical for developing more robust, data-driven models of adaptive behavior.

In outlining numerous challenges, this dissertation may have impressed upon the reader that more questions have been raised than answered. However, this work has made significant progress by introducing a cohesive framework that articulates how environmental characteristics influence adaptive agent behavior. It establishes connections between observable quantities, such as resources and growth, and cognitive and behavioral mechanisms, such as learning and group formation, all grounded in a modular theoretical foundation. This unified quantitative theory will enable theoretically robust social systems research across various scales, inspiring optimism for the future of this emerging field.

APPENDIX A

APPENDIX

A.1 Heterogeneous Growth

This appendix and the following appendices contain analytical derivations for the results presented in this dissertation.

A.1.1 Parameter Covariances

We compute the mean population effective growth rate over a population of size N by separating the expected growth rate into a mean and covariance term

$$\begin{aligned} \langle \bar{\eta} r \rangle_N &= \frac{1}{N} \sum_{j=1}^{N} \bar{\eta}_j r_j, && \text{(A.1)} \\ &= \frac{1}{N} \sum_{j=1}^{N} (\bar{\eta}_j - G) r_j + \frac{1}{N} \sum_{j=1}^{N} G r_j, && \text{(A.2)} \\ &= G\mu_0 + \mathrm{cov}_N(\bar{\eta}_i, r_i), && \text{(A.3)} \\ & && \text{(A.4)} \end{aligned}$$

where the population growth rate can be factored out as $G' = G + \mathrm{cov}_N(\bar{\eta}_i, \frac{r_i}{\mu_0})$.

Assuming r is a lognormal distributed quantity, then so is r/μ_0. Define the normally distributed variable $y \equiv \ln r/\mu_0 \sim \mathcal{N}(0, \sigma_r)$, then the covariance is equal to $\mathrm{cov}_N(\bar{\eta}, e^y)$. The first moment is calculated

$$\langle \bar{\eta} e^y \rangle_N = \int d\bar{\eta} dy \bar{\eta} e^y P(\bar{\eta}, y), \tag{A.5}$$

$$= \int dy e^y \langle \bar{\eta} \rangle_{Y=y} = \int dy e^y (G + \rho \frac{\sigma_{\bar{\eta}}}{\sigma_y} y), \tag{A.6}$$

$$= G\langle e^y \rangle_N + \rho \frac{\sigma_{\bar{\eta}}}{\sigma_y} \langle y e^y \rangle_N = (G + \rho \sigma_{\bar{\eta}} \sigma_y) e^{\sigma_y^2/2}. \tag{A.7}$$

Where the expectation value of ye^y is calculated in [237]. With an expectation value defined, the covariance becomes

$$\mathrm{cov}(\bar{\eta}, e^y) = \langle \bar{\eta} e^{\bar{\eta}} \rangle_N - \langle \bar{\eta} \rangle_N \langle e^y \rangle_N, \tag{A.8}$$

$$= (G + \rho \sigma_{\bar{\eta}} \sigma_y) e^{\sigma_y^2/2} - G e^{\sigma_y^2/2}, \tag{A.9}$$

$$= \rho \sigma_{\bar{\eta}} \sigma_y e^{\sigma_y^2/2}. \tag{A.10}$$

In the main text, I define σ_y as the variance in lognormal resources, σ_r.

A.2 Learning in Shared Environments

A.2.1 Stochastic Growth Model

Multiplying and dividing by $P(e|s)$ in the logarithm of Eq. 2 yields

$$\begin{aligned} \gamma &= \sum_{e,s} P(e,s) \log \left[w_e P(e|s) \frac{X(e|s)}{P(e|s)} \right] \\ &= \sum_{e,s} P(e,s) \log \frac{P(e|s)}{P(e)} - P(s)P(e|s) \log \frac{P(e|s)}{X(e|s)} \\ &= I(E;S) - \mathrm{E}_s\big(D_{KL}\big[P(E|s)||X(E|s)\big]\big), \end{aligned} \tag{A.11}$$

where E_s is an expectation value over all signal states.

Growth rate in the multinomial model

Consider a conditional probability that is degenerate off-diagonal,

$$P(e|s) = f(p,l) = \begin{cases} p & \text{if } s = e, \\ \frac{1-p}{l-1} & \text{if } s \neq e. \end{cases} \tag{A.12}$$

The "correct" outcome corresponding to the sampled event occurs with conditional probability $0 < p \leq 1$, and all other "incorrect" guesses occur with some uniform probability normalized to

$$\sum_{e}^{l-1} P(e|s) = 1 - p; \quad s \neq e. \tag{A.13}$$

We describe the agent's posterior for all agents with the same form, with the "correct" binomial coefficient x. Thus, I calculate the growth rate by taking the expectation value of the posterior over the set of signals, summing over diagonal and off-diagonal components separately.

The mutual information separates into a term of only $l = 1/P(e)$, an on-diagonal, and off-diagonal term

$$\begin{aligned} I(E;S) &= \sum_{e,s}^{l} P(e,s)\big[\log l + \log P(e|s)\big] \\ &= \log l + p\log p + (1-p)\log\frac{1-p}{l-1} \\ &= H(E) - H(E|S), \end{aligned} \tag{A.14}$$

with the entropy of the outcome given by $H(E) = \log l$ and the reduction in entropy by the signal given by $H(E|S) = -p\log p - (1-p)\log\frac{1-p}{l-1}$. The information maximizes as $p \to 1$

and increases with l, and vanishes at $p \to 1/l$. The divergence is

$$\begin{aligned} \mathrm{E}_s\big[D_{KL}(P||X)\big] &= \sum_{e,s} P(e,s) \log \frac{P(e|s)}{X(e|s)} \\ &= p \log \frac{p}{x} + (1-p) \log \frac{1-p}{1-x}, \end{aligned} \tag{A.15}$$

which is always non-negative and vanishes when $x \to p$. We can write the growth rate as the difference between these two terms as

$$\gamma = \mathrm{E}\big[\log lf(x,l)\big] = \log l + p \log x + (1-p) \log \frac{1-x}{l-1}. \tag{A.16}$$

Variance of multinomial growth model

The volatility can be calculated via the second moment of the stochastic growth rate as

$$\sigma = \sqrt{E\big[\log\big(lf(x,l)\big)^2\big] - E\big[\log lf(x,l)\big]^2}. \tag{A.17}$$

$\mathrm{E}\big[\log lf(x,l)\big]$ is simply γ, and the second term is

$$\begin{aligned} \mathrm{E}\big[\log lf(x,l)\big]^2 &= \left(\log l + p \log x + (1-p) \log \frac{1-x}{l-1}\right)^2 \\ &= \log^2 l + p^2 \log^2 x + (1-p)^2 \log^2 \frac{1-x}{l-1} \\ &+ 2p \log l \log x + 2(1-p) \log l \log \frac{1-x}{l-1} \\ &+ 2p(1-p) \log x \log \frac{1-x}{l-1}. \end{aligned} \tag{A.18}$$

The first term expands to

$$\begin{aligned}\mathrm{E}\big[\log\big(lf(x,l)\big)^2\big] &= \mathrm{E}_{e,s}\big[(\log P(e|s)+\log l\big)^2\big] \\ &= \log^2 l + p\log^2 x + (1-p)\log^2\frac{1-x}{l-1} \\ &+ 2p\log l\log x + 2(1-p)\log l\log\frac{1-x}{l-1}.\end{aligned} \tag{A.19}$$

Combining these two quantities yields the volatility, where $(1-p)-(1-p)^2 = p(1-p)$,

$$\begin{aligned}\sigma_n &= \sqrt{p(1-p)\left[\log^2 x + \log^2\frac{1-x}{l-1} - 2\log x\log\frac{1-x}{l-1}\right]} \\ &= \sqrt{p(1-p)}\log\frac{x(l-1)}{1-x}.\end{aligned} \tag{A.20}$$

The variance of investment clusters of size $1/\omega$ scales as

$$\sigma_t^2 = \frac{1}{\gamma}\sigma_n^2, \tag{A.21}$$

where the subscript t denotes the temporal variance.

A.2.2 Bayesian Inference with Latent Dirichlet Allocation

NOTE: UPDATE FROM UPDATED SUPPLEMENT, TYPOS BELOW

In this section, I derive the Latent Dirichlet Allocation (LDA) mode for the degenerate multinomial environment. The Bayesian update equation is given by

$$X(e|s) \propto \frac{\big(m_{(-s)}^{(-e)} + \tilde{\beta}_s^e\big)}{\big(M^{(-s)} + \tilde{B}^s\big)}(n_{(-e)} + \tilde{\alpha}_e), \tag{A.22}$$

for $m_{(-s)}^{(-e)}$ number of samples of outcome s conditional on e excluding the current, $n_{(-e)}$

the number of samples of e excluding the current in a batch of $n = \sum_e n_{(-e)}$ trials, where $M^{(-s)} = \sum_e m^{(-e)}_{(-s)}$. We set $\alpha_e = 1$, as every event is equally likely. For $s = e$, $\tilde{\beta}^e_s = x_e$, and for $s \neq e$, $\beta_{es} = \frac{(l-1)}{1-x}$ to impose degenerate off-diagonal conditions on $s|e$. We introduce $\tilde{B}^s = \sum_e \tilde{\beta}^e_s$, whereby symmetry, $\tilde{B}^s \equiv \tilde{B} = 1$, and I count over the diagonals, $n_{e=s}$, and off diagonals, $n_{e\neq s}$. Therefore the diagonal environmental posterior is

$$P(e|s) \propto \frac{\left(m^{(-e)}_{(-s=e)} + x_e\right)}{\left(M^{(-s)} + 1\right)}(n_{(-e)} + 1), \tag{A.23}$$

and the off-diagonal is

$$P(e|s) \propto \frac{\left(m^{(-e)}_{(-s\neq e)} + \frac{1-x_e}{l-1}\right)}{\left(M^{(-s)} + 1\right)}(n_{(-e)} + 1). \tag{A.24}$$

Asymptotic, temporal behavior

We introduce the temporal behavior, with two constants. We multiply the number of observations by the observation rate ω, with units *samples/time* and the inference rate k, with unit *time/update*. The inference rate counts the number of samples per Bayesian update, and the observation rate counts the updates per unit time. We introduce the inference time, k, a hyperprior magnitude that weighs the evidence versus the prior, leaving

$$P(e|s) \propto \frac{(m^{(-e)}_{(-s)}/\omega + \tilde{\beta}^e_s k)}{(M^{(-s)}/\omega + 1k)}(n_{(-e)}/\omega + \tilde{1}/k). \tag{A.25}$$

Over many observations, the law of large numbers argues that each outcome count converges to the environmental posterior with some noise, ξ_i as

$$
\begin{aligned}
M^{(-s)}/\omega &\to P(s)Nt + \xi_s \\
n_{(-e)}/\omega &\to P(e)Nt + \xi_e \\
m^{(-e)}_{(-s)}/\omega &\to P(s|e)Nt + \xi_{s|e},
\end{aligned}
\tag{A.26}
$$

where the ξ's are fluctuation terms representing deviations from the mean. Over many i.i.d observations of events, $\xi \to 0$. The marginal terms converge to uniform over all states, and the agent posterior converges to the dynamical distribution

$$
X(e, \lambda|s) = \frac{P(s|e)\lambda + X(s|e)}{1 + \lambda},
\tag{A.27}
$$

where I have converted to the time domain $t = N/\omega$, and substituted the dimensionless inference sample size $\lambda = t/kl$. Over long times, the distribution converges to the environmental posterior by

$$
\begin{aligned}
X(e, \lambda|s) &\propto \frac{P(s|e)\lambda + X(s|e)}{P(s)\lambda + 1}\big(P(e)\lambda + \alpha_e\big) \\
&\to \left(P(s|e) + \frac{X(s|e, 0)}{\lambda}\right)\frac{P(e)}{P(s)} = P(e|s),
\end{aligned}
\tag{A.28}
$$

yielding power law time-averaged behavior. At early times, as $t \to 0$ the posterior is proportional to the agent's initial agent posterior, $X(E|S)$, and converges to $P(E|S)$ as $kl \ll t \to \infty$. If agents are initialized with the same diagonal posterior value such that $X(s|e) = X(e', s')$ for all $e = e, s' = s'$, We can assume that the diagonals of an agent uniformly converge to p in time such that $X(s|e) \propto x(t)$ for all $s = e$,

A.2.3 Population variance of the average growth rate

The mean growth rate is computed, where for brevity, the expected divergence for agent i with signals $s_i \in S_i$ is given as $\mathrm{E}_{s_i}\big(D_{KL}\big[P(E|s_i)||X(E|s_i\big]\big) \equiv D_i$, and the mutual information between individual signals and the environment, $I(E;S_i) \equiv I_i$

$$\begin{aligned}\langle \gamma_i \rangle &= \frac{1}{N}\sum_i I(E;S_i) - \mathrm{E}_{s_i}\big(D_{KL}\big[P(E|s_i)||X(E|s_i\big]\big) \\ &= \big\langle I_i \big\rangle - \langle D_i \rangle,\end{aligned} \tag{A.29}$$

where angle brackets denote population arithmetic means. The variance in growth rates is calculated

$$\begin{aligned}\mathrm{Var}_N[\gamma_i] &= \big\langle (\gamma_i - \langle \gamma_i \rangle)^2 \big\rangle, \\ &= \big\langle \gamma_i^2 + \langle \gamma_i \rangle^2 - 2\gamma_i \langle \gamma_i \rangle \big\rangle \\ &= \langle I_i^2 \rangle - \langle I_i \rangle^2 + \langle D_i^2 \rangle - \langle D_i \rangle^2 \\ &- 2\big(\langle I_i D_i \rangle - \langle I_i \rangle \langle D_i \rangle\big) \\ &= \mathrm{Var}_N[I_i] + \mathrm{Var}_N[D_i] - 2\mathrm{Covar}_N\big[I_i D_i\big].\end{aligned} \tag{A.30}$$

When all agents are exposed to the same environment and share the same likelihood, the first and third terms vanish, leaving

$$\mathrm{Var}_N[\gamma_i] = \mathrm{Var}_N\bigg[\mathrm{E}_{s_i}\bigg(D_{KL}\big[P(S|s_i)||X(E|s_i)\big]\bigg)\bigg]. \tag{A.31}$$

Multinomial parameter variance

The binomial variance can be computed exactly as

$$\begin{aligned}\mathrm{Var}_N\big[x_i(\lambda)\big] &= \frac{1}{N}\sum_i^N \left[\frac{p\lambda + x_i}{1+\lambda}\right]^2 - \left[\frac{p\lambda + \langle x_j\rangle}{1+\lambda}\right]^2 \\ &= \frac{1}{N}\sum_i^N \left[2\frac{x_i p\lambda + x_i^2}{(1+\lambda)^2} - 2\frac{\langle x_j\rangle p\lambda - \langle x_j\rangle^2}{(1+\lambda)^2}\right] \\ &= \frac{\langle x^2\rangle - \langle x_j\rangle^2}{(1+\lambda)^2} = \frac{\sigma_x^2}{(1+\lambda)^2}.\end{aligned} \tag{A.32}$$

Multinomial growth rate variance

The variance of a function, $\gamma(x)$, of a random variable, x, is given generally by the Taylor expansion of that function [62]. It is written as

$$\begin{aligned}\mathrm{Var}_N\big(\gamma[x_i(\lambda)]\big) = \gamma'\big[\big\langle x_i(\lambda)\big\rangle\big]\mathrm{Var}_N\big[x_i(\lambda)\big] \\ - \frac{\gamma''\big[\big\langle x_i(\lambda)\big\rangle\big]^2}{4}\mathrm{Var}_N^2\big[x_i(\lambda)\big] + \bar{T}^3,\end{aligned} \tag{A.33}$$

where primes denote differentiation with respect to x, and $\bar{T}^3$ are higher order terms that are only relevant at small times. The first and second-order derivatives of γ are given by

$$\begin{aligned}\gamma'(x) &= \frac{p}{x} - \frac{1-p}{1-x} \\ \gamma''(x) &= -\left[\frac{p}{x^2} + \frac{1-p}{(1-x)^2}\right],\end{aligned} \tag{A.34}$$

and the variance term is given by

$$\mathrm{Var}_N\big[x_i(\lambda)\big] = \frac{\sigma_x^2}{(1+\lambda)^2}. \tag{A.35}$$

The growth rate variance after small times is given by

$$\begin{aligned}\mathrm{Var}_N\big(\gamma[x_i(\lambda)]\big) = &\left[\frac{p}{\bar{x}} - \frac{1-p}{1-\bar{x}}\right]\frac{\sigma_x^2}{(1+\lambda)^2} \\ &+ \left[\frac{p}{\bar{x}^2} + \frac{1-p}{(1-\bar{x})^2}\right]\left[\frac{\sigma_x^2}{(1+\lambda)^2}\right]^2,\end{aligned} \tag{A.36}$$

where for brevity, $\bar{x} \equiv \big\langle x_i(\lambda)\big\rangle$.

A.3 Information Sharing

A.3.1 Derivation of Information Synergy

Consider a target statistical variable E (environment), that I wish to predict using l other variables (signals) $\boldsymbol{S} = \{S_1, \ldots, S_l\}$. The mutual information between each signal S_i separately and E is given by [25]

$$I(E; S_i) = H(E) - H(E|S_i) = -\frac{\Delta H(E)}{\Delta S_i}. \tag{A.37}$$

where $H(E)$ is the Shannon entropy of E, and the variation measures the difference in entropy of the event when conditioned on the signal. From the rules of information aggregation, this expression generalizes to information across every added signal [32]. The mutual information between the event and the set of several signals is given by

$$I(E; \boldsymbol{S}) = -\sum_{i=1}^{l}\frac{\Delta H(E)}{\Delta S_i} - \sum_{i>j=1}^{l}\frac{\Delta^2 H(E)}{\Delta S_i \Delta S_j} - \cdots - \frac{\Delta^l H(E)}{\Delta S_1 \ldots \Delta S_l}. \tag{A.38}$$

The first term of this expansion is just a sum over the mutual information of each individual signal and the environment. The goal of this section is to show that the inclusion of each new signal introduces a coefficient of redundancy of progressively higher order. The first term is

$$\begin{aligned}\frac{\Delta^2 H(E)}{\Delta S_i \Delta S_j} =& H(S_i) - H(S_i|E) + H(S_j) - H(S_j|E) - H(S_i, S_j) + H(S_i, S_j|E) \\ =& H(S_i) - H(S_i|E) - H(S_i|S_2) + H(S_i|S_j, E) \\ =& I(S_i; S_j) - I(S_i; S_j|E) \equiv R(E; S_i; S_j),\end{aligned} \tag{A.39}$$

where I used the identity $H(A, B) = H(A|B) + H(B)$. We denote R as the coefficient of redundancy, which measures the difference in mutual information between the variables, $I(\boldsymbol{S}) \equiv I(S_1; \ldots, S_k)$, and the mutual information of the variables conditioned on E, $I(\boldsymbol{S}|E)$. When $I(S_i; S_j) < I(S_i; S_j|R)$, the signals contain less mutual information in the absence of the event (we gain information by considering the event), and $R(S_i; S_j; X) < 0$. In this case agents experience a positive benefit from pooling information, which I call synergy.

To demonstrate this effect to higher orders in goups of signals, I perform a similar calculation for a three-signal interaction.

$$\begin{aligned}\frac{\Delta^3 H(E)}{\Delta S_i \Delta S_j \Delta S_k} =& H(E) - H(E|\{S_i, S_j, S_k\}) \\ =& H(S_i) + H(S_j) + H(S_k) - H(S_i|E) - H(S_j|E) - H(S_k|E) - H(S_i, S_j) \\ & - H(S_j, S_k) - H(S_k, S_i) + H(S_i, S_j|E) + H(S_j, S_k|E) + H(S_k, S_i|E) \\ & + H(S_i, S_j, S_k) - H(S_i, S_j, S_k|E) \\ =& H(S_i, S_j, S_k) - H(S_i|S_j) - H(S_j|S_k) - H(S_k|S_i) + H(S_i|S_j, E) \\ & + H(S_j|S_k, E) + H(S_k|S_i, E) - H(S_i, S_j, S_k|E) \\ =& I(S_i; S_j; S_k) - I(S_i; S_j; S_k|E) \equiv R(E; S_i, S_j, S_k).\end{aligned} \tag{A.40}$$

We see that an analogous redundancy coefficient arises in three dimensions. This can generally be retrieved for arbitrary number of dimensions through a similar iterative procedure.

I refer to the sum of these moments collectively as the redundancy of the joint distribution, denoted R_P [32],

$$\begin{aligned} R_P &\equiv -\sum_{i>j=1}^{l} \frac{\Delta^2 H(E)}{\Delta S_i \Delta S_j} - \cdots - \frac{\Delta^l H(E)}{\Delta S_1 \ldots \Delta S_l} \\ &= \sum_{i>j=0}^{l} \big[I(S_i; S_j) - I(S_i; S_j|E) \big] + \cdots + I(S_1; \ldots; S_l) - I(S_1, \ldots, S_l|E) \end{aligned} \quad . \tag{A.41}$$

Note that redundancies of lower order than cardinality of the signal space must be computed over every combination of signals. For example, when $l = 3$, there are three second order redundancy terms.

This expansion generally defines the benefits to cooperation over increasingly higher orders of cooperation (number of signals). This expression can be used to compute the relative strengths of the various orders of interaction for any set of signals and environmental variables, given their conditional distributions.

Stochastic Growth in Synergistic Environments

Consider an environment with events conditionally dependent on signals characterized by a joint distribution $P(E, \boldsymbol{S})$ for event E and l signals S. Consider a cooperative Kelly investment scheme whereby each participant, agent i, witnesses signal $s_i \in S_i$, and informs the collective how to invest their shared resources r. The mechanics of pooling resources and collectively investing will be discussed below. Kelly's formalism can be adapted by expanding the environmental probability to contain l signals, $P(E, S) \rightarrow P(E, \boldsymbol{S})$, as can the betting matrix $X(E|S) \rightarrow X(E|\boldsymbol{S})$, where $\boldsymbol{S} = \{S_1, \ldots, S_l\}$. When odds are fair, the Kelly growth rate is given by the returns to each investment, averaged over the probability of that signal, event pair

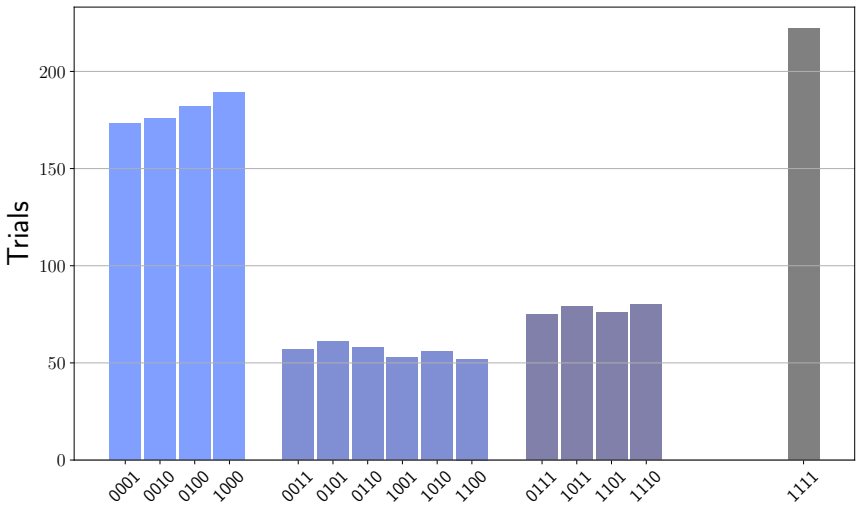

Figure A.1: Trial count by group identification for N=5000. Zeros denote the exclusion of a signal, and ones the inclusion.

$$G = \sum_{e,\boldsymbol{s}}^{E,\boldsymbol{S}} p(e,\boldsymbol{s}) \log \frac{x(e|\boldsymbol{s})}{p(e)}. \tag{A.42}$$

We expand this equation by inserting $p(e,\boldsymbol{s})$ into the numerator and denominator of the log

$$G = \sum_{e,\boldsymbol{s}}^{E,\boldsymbol{s}} p(e,\boldsymbol{s}) \left[\log \frac{p(e,\boldsymbol{s})}{p(\boldsymbol{s})p(e)} - \log \frac{p(e|\boldsymbol{s})}{x(e|\boldsymbol{s})} \right]. \tag{A.43}$$

These two terms can be simply expressed as $G = I(E;\boldsymbol{S}) - \mathrm{E}_{\boldsymbol{s}}\big[D_{KL}\big(P(E|\boldsymbol{s})||X(E|\boldsymbol{s})\big)\big]$, similar to previous work, but we can decomposes this equation in terms of redundant information across the signals using equations (A.38) and (A.41).

A.3.2 Monte-Carlo Simulation Details

To study the dynamics of resources in the UXOR environment, I simulated agent investments in a Monte Carlo sampled environment. We randomly assigned $N = 5000$ agents signals in an $l = 4$ environment, then randomly assigned them to groups sized $l \leq N_g \leq 11$. This

results in an ensemble of groups with cooperants $1 \leq k_g \leq 4$. Each instance of a group represents one sample of that group configuration. This random assignment scheme, selected for convenience, results in a different sample size for each group identity, where identity is given by the types of signals present. The sample size of a group type scales inversely with the number of combinations of signals of each order l. However, each group combination is sampled between 50 and 225 times, resulting in resolvable statistical behavior across all group types. Figure S1 gives the statistics of the groups simulated in Figure 4 of the main text.

The entire population, divided into groups, is instantiated in an environment with the same environmental parameters. The UXOR gate is defined with parameters: $p_1 = .1473$, $p_2 = .7236$, $p_3 = .46451306$, $p_4 = .7068$, $p_5 = .4658$, $p_6 = .3594$, $p_7 = .9708$, $p_8 = .3709$, $p_9 = .8788$, $p_{10} = .7094$, $p_{11} = .7972$, $p_{12} = .6380$, $p_{13} = .4327$, $p_{14} = .2238$, $p_{15} = .4044$, $p_{16} = 0.15949486$. These parameters were discovered from a random search of parameter space and were selected for their convenient qualitative properties.

During each step of the simulation, I reveal Bernoulli-sampled signals to the groups, whose agents make collective decisions on which events to allocate resources. In this sense, every agent in the group views the same signals, and the group decides the subspace of signals available via the members. Allocations and resource and information rewards are processed identically to prior works, where an LDA defines the update to a group's decision parameters. For each group, I track the resources of a representative agent, informed by the group, investing their individual resources through time.

A.4 Mutual Adaptation

A.4.1 Defining the Deterministic ODE

This section was developed in collaboration with Prof. Max-Olivier Hongler of the EPFL School of Engineering.

The equation for the mean of the Gaussian Gamma for variable x is given by

$$\mu_x^n = \frac{\sum_{i=0}^{n} y_i + \mu_0\alpha}{n+1+\alpha} = \frac{y_{n-1} + \sum_{i=0}^{n-1} y_i + \mu_0\alpha}{n+1+\alpha}$$

The sample y_n corresponds to the mean of $P(y)$, μ_y, plus some noise ξ_n. In the deterministic case, $\xi_n = 0$, and remaining two terms constitute $\mu_x^{n-1}(n+\alpha)$. We can therefore redefine this quantity as Eq. 5.2 from the main text

$$\mu_x^n = \frac{\mu_x^{n-1}\big[(n-1)/\omega + \alpha\big] + \mu_y^{n-1}}{\frac{n}{\omega} + \alpha}, \tag{A.44}$$

where I have written $n \to \frac{n}{\omega}$ to enable a conversion to continuous time variables later on. In the deterministic case, I define the difference operator

$$\begin{aligned}\Delta\mu_x = \mu_x^n - \mu_x^{n-1} &= \frac{\mu_x^{n-1}\big[(n-1)/\omega + \alpha\big] + \mu_y^{n-1}}{\frac{n}{\omega} + \alpha} - \frac{\mu_x^{n-1}(\frac{n}{\omega} + \alpha)}{\frac{n}{\omega} + \alpha} \\ &= \frac{\mu_y^n - \mu_x^n}{\frac{n}{\omega} + \alpha}\end{aligned} \tag{A.45}$$

which as $\omega \to \infty$ converges to the continuous time ODE used in the main text. Finally, Gaussian noise will be linearly added to the deterministic dynamics.

Deterministic evolution of preferences

I start by solving the deterministic motion, namely the set of ODEs:

$$\begin{cases} \frac{d\mu_x(t)}{dt} = \frac{\mu_y(t)-\mu_x(t)}{t+\alpha}, & \mu_x(0) = x_0 \\ \\ \frac{d\mu_y(t)}{dt} = \frac{\mu_x(t)-\mu_y(t)}{t+\beta}, & \mu_y(0) = y_0. \end{cases} \tag{A.46}$$

To process further, I introduce the new variables:

$$\mu_\Delta(t) = \left[\mu_x(t) - \mu_y(t)\right] \qquad \text{and} \qquad \mu_\Sigma(t) = \left[\mu_x(t) + \mu_y(t)\right]. \tag{A.47}$$

In terms of the new variables, Eq. (A.46) reads:

$$\begin{cases} \left[(t+\alpha)(t+\beta)\right] \frac{d\mu_\Delta(t)}{dt} = -\left[2t+\alpha+\beta\right] \mu_\Delta(t), \\ \\ \left[(t+\alpha)(t+\beta)\right] \frac{d\mu_\Sigma(t)}{dt} = \left[\alpha-\beta\right] \mu_\Delta(t). \end{cases} \tag{A.48}$$

From Eq. (A.48), I immediately have:

$$\begin{cases} \frac{d\ln(\mu_\Delta(t))}{dt} = -\frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]} = -\frac{d\ln[(t+\alpha)(t+\beta)]}{dt} \;\Rightarrow\; \mu_\Delta(t) = \frac{\Delta_0\alpha\beta}{[(t+\alpha)(t+\beta)]} := \frac{\Delta_0\alpha\beta}{R(t)}, \\ \\ \left[(t+\alpha)(t+\beta)\right] \frac{d\mu_\Sigma(t)}{dt} = \left[\alpha-\beta\right] \mu_\Delta(t) \quad \Rightarrow \quad \frac{d\mu_\Sigma(t)}{dt} = \frac{(\alpha-\beta)\alpha\beta\Delta_0}{[(t+\alpha)(t+\beta)]^2} = \Delta_0 \frac{\alpha\beta(\alpha-\beta)}{R^2(t)} \end{cases} \tag{A.49}$$

with:

$$R(t) := \left[(t+\alpha)(t+\beta)\right] \tag{A.50}$$

By direct calculation, one may verify the identities:

$\int \frac{1}{R(t)}dt = \frac{1}{\alpha-\beta}\ln\left[\frac{t+\beta}{t+\alpha}\right]$

$$\int \frac{1}{R^2(t)}dt = -\frac{1}{(\alpha-\beta)^2}\left\{\frac{d\ln[R(t)]}{dt} + 2\int \frac{1}{R(t)}dt\right\} = -\frac{1}{(\alpha-\beta)^2}\left\{\frac{d\ln[R(t)]}{dt} + \frac{2}{\alpha-\beta}\ln\left[\frac{t+\beta}{t+\alpha}\right]\right\} =$$

$$-\frac{1}{(\alpha-\beta)^2}\left\{\frac{2t+\alpha+\beta}{t^2+(\alpha+\beta)t+\alpha\beta} + \frac{2}{\alpha-\beta}\ln\left[\frac{t+\beta}{t+\alpha}\right]\right\}. \tag{A.51}$$

Accordingly, from Eqs. (A.49) and (A.51), I obtain:

$$\mu_\Sigma(t) = \int \frac{(\alpha-\beta)\alpha\beta\Delta_0}{R^2(t)}dt = -\alpha\beta\Delta_0\left\{\frac{2t+\alpha+\beta}{(\alpha-\beta)R(t)} + \frac{2}{(\alpha-\beta)^2}\ln\left[\frac{t+\beta}{t+\alpha}\right]\right\} + C \tag{A.52}$$

where C is an integration constant to be determined by the initial condition. At time $t = 0$, I have:

$$C = \Sigma_0 + \Delta_0\left\{\frac{\alpha+\beta}{(\alpha-\beta)} + \frac{2\alpha\beta[\ln\beta - \ln\alpha]}{(\alpha-\beta)^2}\right\}. \tag{A.53}$$

Hence we can write:

$$\mu_\Sigma(t) = \Sigma_0 + \frac{\Delta_0}{(\alpha-\beta)}\left[\alpha+\beta - \frac{(2t+\alpha+\beta)\alpha\beta}{(t+\alpha)(t+\beta)} + \frac{2\alpha\beta}{(\alpha-\beta)}\ln\left(\frac{\beta t+\alpha\beta}{\alpha t+\alpha\beta}\right)\right]. \tag{A.54}$$

Note that I have:

$$\begin{cases} \lim\limits_{t\to 0^+} \mu_\Sigma(t) = \Sigma_0, \\ \\ \lim\limits_{t\to\infty} \mu_\Sigma(t) := \mu_{\Sigma,f} = \Sigma_0 + \Delta_0 \left[\dfrac{\alpha+\beta}{\alpha-\beta}\right] = 2\left[\dfrac{\alpha x_0 - \beta y_0}{\alpha-\beta}\right] + 2\Delta_0 \dfrac{\alpha\beta \ln[\beta/\alpha]}{(\alpha-\beta)^2}, \quad \alpha \neq \beta, \\ \\ \lim\limits_{t\to\infty} \mu_\Sigma(t) := \Sigma_0, \quad \alpha = \beta, \end{cases} \tag{A.55}$$

In terms of $\mu_x(t)$ and $\mu_y(t)$, I have:

$$\begin{cases} \mu_x(t) = \frac{\mu_\Delta(t)+\mu_\Sigma(t)}{2} = \left[\frac{\alpha x_0 - \beta y_0}{\alpha-\beta}\right] + (x_0 - y_0)\left\{\frac{\alpha\beta}{\alpha-\beta}\ln\left[\frac{\alpha\beta+\beta t}{\alpha\beta+\alpha t}\right] - \frac{\alpha\beta}{(\alpha-\beta)(t+\alpha)}\right\}. \\ \\ \mu_y(t) = \frac{\mu_\Delta(t)-\mu_\Sigma(t)}{2} = \left[\frac{\alpha x_0 - \beta y_0}{\alpha-\beta}\right] + (x_0 - y_0)\left\{\frac{\alpha\beta}{\alpha-\beta}\ln\left[\frac{\alpha\beta+\beta t}{\alpha\beta+\alpha t}\right] - \frac{\alpha\beta}{(\alpha-\beta)(t+\beta)}\right\} \end{cases} \tag{A.56}$$

which can be summarised as:

$$\begin{pmatrix} \mu_x(t) \\ \\ \mu_y(t) \end{pmatrix} = \mathbb{M}(t) \begin{pmatrix} x_0 \\ \\ y_0 \end{pmatrix}, \tag{A.57}$$

where the matrix $\mathbb{M}(t)$ reads as:

$$\mathbb{M} = \begin{pmatrix} -M_1(0) + M_2(t) & M_2(0) - M_2(t) \\ -M_1(0) + M_1(t) & M_2(0) - M_1(t) \end{pmatrix},$$

$$M_1(t) = \frac{\alpha\beta}{\alpha-\beta} \ln\left[\frac{\alpha\beta+\beta t}{\alpha\beta+\alpha t}\right] - \frac{\alpha\beta}{(\alpha-\beta)(t+\alpha)}, \tag{A.58}$$

$$M_2(t) = \frac{\alpha\beta}{\alpha-\beta} \ln\left[\frac{\alpha\beta+\beta t}{\alpha\beta+\alpha t}\right] - \frac{\alpha\beta}{(\alpha-\beta)(t+\beta)}.$$

Stationary regime

From Eqs. (A.57) and (A.58), one immediately concludes that the final state $(\mu_x(\infty), \mu_x(\infty))$ is given by:

$$\begin{pmatrix} \mu_x(\infty) \\ \\ \mu_y(\infty) \end{pmatrix} = \begin{pmatrix} \frac{\alpha}{\alpha-\beta} & \frac{\beta}{\beta-\alpha} \\ \\ \frac{\alpha}{\alpha-\beta} & \frac{\beta}{\beta-\alpha} \end{pmatrix} \begin{pmatrix} x_0 \\ \\ y_0 \end{pmatrix} := \mathbb{M}_\infty \begin{pmatrix} x_0 \\ \\ y_0 \end{pmatrix} \tag{A.59}$$

We observe that $\big(\mu_x(\infty), \mu_y(\infty)\big)$ depends on the initial condition (x_0, x_0). Note that $\alpha = \beta$ is a singular situation. In addition observe also that for $\alpha = 0$, I obtain $\mu_x(\infty) = \mu_y(\infty) = y_0$ and conversely for $\beta = 0$, I have $\mu_x(\infty) = \mu_y(\infty) = x_0$ thus showing that in both of these limiting cases, the evolution affects a single variable.

A.4.2 The (X_t, Y_t) stochastic process using coupled dynamics

Consider the stochastic process $(X_t, Y_t) \in \mathbb{R}^2$:

$$\begin{cases} dX_t = \frac{-\Delta_t dt + \sigma_y dW_{1,t}}{t+\alpha} = \frac{(y-x)dt + \sigma_y dW_{1,t}}{t+\alpha}, \qquad X_0 = x_0 \\ \\ dY_t = \frac{+\Delta_t dt + \sigma_x dW_{2,t}}{t+\beta} = \frac{(x-y)dt + \sigma_x dW_{2,t}}{t+\beta}, \qquad Y_0 = y_0. \end{cases} \tag{A.60}$$

where $dW_{1,t}$ and $dW_{2,t}$ are independent Gaussian Noise processes (WGN). To study the

(X_t, Y_t) bi-variate Gaussian [1] and Markovian diffusion process, it is advantageous to proceed with the change of variables:

$$\begin{cases} \begin{pmatrix} X_t \\ Y_t \end{pmatrix} \mapsto \begin{pmatrix} \Delta_t := X_t - Y_t \\ \Sigma_t := X_t + Y_t \end{pmatrix}, \\ \\ d\Delta_t = -\frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]}\Delta_t dt + \frac{\sqrt{[\sigma_x^2(t+\beta)^2+\sigma_y^2(t+\alpha)^2]}dB_{1,t}}{[(t+\alpha)(t+\beta)]}, \\ \\ d\Sigma_t = +\frac{(\alpha-\beta)}{[(t+\alpha)(t+\beta)]}\Delta_t dt + \frac{\sqrt{[\sigma_x^2(t+\beta)^2+\sigma_y^2(t+\alpha)^2]}dB_{2,t}}{[(t+\alpha)(t+\beta)]}, \end{cases} \tag{A.61}$$

where I have used the property:

$$\begin{cases} \frac{\sigma_y}{(t+\alpha)}dW_{1,t} - \frac{\sigma_x}{(t+\alpha)}dW_{2,t} = \frac{\sqrt{[\sigma_x^2(t+\beta)^2+\sigma_y^2(t+\alpha)^2]}dB_{1,t}}{[(t+\alpha)(t+\beta)]}, \\ \\ \frac{\sigma_y}{(t+\alpha)}dW_{1,t} + \frac{\sigma_x}{(t+\alpha)}dW_{2,t} = \frac{\sqrt{[\sigma_x^2(t+\beta)^2+\sigma_y^2(t+\alpha)^2]}dB_{2,t}}{[(t+\alpha)(t+\beta)]} \end{cases} \tag{A.62}$$

with $dB_{1,t}$ and $dB_{2,t}$ being now **correlated** WGN's. Since the Δ_t process is actually decoupled from the Σ_t, I shall proceed in two steps.

A.4.3 D.2.1 The Δ_t process

The probabilistic properties of Δ_t stochastic process in Eq. (A.61) are fully described by the TPD $P_\Delta(x,t|x_0,0)dx := \text{Prob}\{x \leq \Delta_t \leq (x+dx)|x_0\}$ which solves the FPE:

$$\begin{cases} \partial_t P_\Delta = \partial_x \left\{ \left[\frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]} \right] x \right\} + D(t)\partial_{xx} P_\Delta, \\ \\ D(t) := \frac{\sigma_y^2(t+\beta)^2+\sigma_x^2(t+\alpha)^2}{2(t+\alpha)^2(t+\beta)^2}. \end{cases} \tag{A.63}$$

1. The Gaussian property is ensured since the drifts are linear and linear transformations of Gaussian restitutes Gaussian.

To solve Eq.(A.63), I express the evolution in terms of the constant of the motion Δ_o. Accordingly, I introduce the change of variables:

$$\begin{cases} t \mapsto \tau = t \quad \Rightarrow \quad \partial_t \mapsto \frac{\partial x'}{\partial t}\partial_{x'} + \frac{\partial \tau}{\partial t}\partial_\tau = \left[\frac{2t+\alpha+\beta}{\alpha\beta}\right] x\partial_{x'} + \partial_\tau = \left[\frac{2t+\alpha+\beta}{(t+\alpha)(t+\beta)}\right] x'\partial_{x'} + \partial_t, \\ \\ x \mapsto x' := x\frac{(t+\alpha)(t+\beta)}{\alpha\beta} \quad \Rightarrow \quad \partial_x \mapsto \frac{\partial x'}{\partial x}\partial_{x'} + \frac{\partial \tau}{\partial x}\partial_\tau = \frac{(t+\alpha)(t+\beta)}{\alpha\beta}\partial_{x'} \end{cases} \tag{A.64}$$

In terms of the (t, x'), Eq. (A.63) is transformed into a pure diffusion process. This can be seen as follows, (omitting the arguments of P):

$$\partial_t P_\Delta + \left[\frac{2t+\alpha+\beta}{(t+\alpha)(t+\beta)}\right] x'\partial_{x'} P_\Delta =$$

$$\frac{(\alpha+t)(\beta+t)}{\alpha\beta}\partial_{x'}\left\{\frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]}\frac{\alpha\beta}{(t+\alpha)(t+\beta)}x'P\right\} + D(t)\frac{(\alpha+t)^2(\beta+t)^2}{\alpha^2\beta^2}\partial_{x'x'}P_\Delta =$$

$$\frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]}\partial_{x'}(x'P_\Delta) + D(t)\frac{(\alpha+t)^2(\beta+t)^2}{\alpha^2\beta^2}\partial_{x'x'}P_\Delta,$$

yielding

$$\partial_t P_\Delta = \frac{[2t+\alpha+\beta]}{[(t+\alpha)(t+\beta)]}P_\Delta + D(t)\frac{(\alpha+t)^2(\beta+t)^2}{\alpha^2\beta^2}\partial_{x'x'}P_\Delta. \tag{A.65}$$

Writing $P_\Delta := (t+\alpha)(t+\beta)Q$, Eq(A.65), reduces to the pure time inhomogeneous diffusion:

$$\partial_t Q_\Delta = \left[D(t)\frac{(\alpha+t)(\beta+t)}{\alpha^2\beta^2}\right]\partial_{x'x'}Q_\Delta = \left[\frac{\sigma_x^2(t+\alpha)^2 + \sigma_y^2(t+\beta)^2}{2\alpha^2\beta^2(\alpha+t)(\beta+t)}\right]\partial_{x'x'}Q_\Delta, \tag{A.66}$$

Finally, I introduce the time re-scaling :

$$t \mapsto s(t) := \int_0^t \left[\frac{\sigma_x^2(\alpha+\xi)^2 + \sigma_y^2(\beta+\xi)^2}{\alpha^2\beta^2(\alpha+\xi)(\beta+\xi)} \right] d\xi = \frac{(\sigma_x^2+\sigma_y^2)t + (\alpha-\beta)(\sigma_x^2 \ln\frac{\beta+t}{\beta} - \sigma_y^2 \ln\frac{\alpha+t}{\alpha})}{\alpha^2\beta^2}, \tag{A.67}$$

This enables to rewrite Eq. (A.66) as:

$$\partial_s Q_\Delta = \frac{1}{2}\partial_{yy} Q_\Delta \qquad \Rightarrow \qquad Q_\Delta = \frac{1}{\sqrt{2\pi s(\tau)}} \exp\left[-\frac{(y-y_0)^2}{2s(\tau)} \right]. \tag{A.68}$$

Proceeding backwards to the nominal (x,t) variables, one ends with:

$$\begin{aligned} lP_\Delta(x,t|x_0,0)dx &= \frac{(t+\alpha)(t+\beta)}{\alpha\beta\sqrt{2\pi s(t)}} \exp\left[-\frac{(x\frac{(t+\alpha)(t+\beta)}{\alpha\beta} - x_0)^2}{2s(\tau)} \right] dx \\ &= \frac{1}{\sqrt{2\pi\sigma_\Delta^2(t)}} \exp\left[-\frac{\big(x - \frac{\alpha\beta\, x_0}{(t+\alpha)(t+\beta)}\big)^2}{2\sigma_\Delta^2(t)} \right], \end{aligned} \tag{A.69}$$

where I used the notation:

$$\begin{cases} \sigma_\Delta^2(t) := \frac{\alpha^2\beta^2 s(t)}{(t+\alpha)^2(t+\beta)^2} = \frac{(\sigma_x^2+\sigma_y^2)t + (\alpha-\beta)\big(\sigma_x^2 \ln\frac{\beta+t}{\beta} - \sigma_y^2 \ln\frac{\alpha+t}{\alpha}\big)}{(t+\alpha)^2(t+\beta)^2}, \\ \\ \lim\limits_{t\to\infty} \sigma_\Delta^2(t) = 0. \end{cases} \tag{A.70}$$

REFERENCES

[1] Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. Technical report, National Bureau of Economic Research, 2017.

[2] Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *The review of economic studies*, 89(1):45–86, 2022.

[3] Daniel Acuna, Paul Schrater, et al. Bayesian modeling of human sequential decision-making on the multi-armed bandit problem. In *Proceedings of the 30th annual conference of the cognitive science society*, volume 100, pages 200–300. Citeseer, 2008.

[4] Ufuk Akcigit and William R Kerr. Growth through heterogeneous innovations. *Journal of Political Economy*, 126(4):1374–1443, 2018.

[5] Paul H Algoet and Thomas M Cover. Asymptotic optimality and asymptotic equipartition properties of log-optimum investment. *The Annals of Probability*, pages 876–898, 1988.

[6] Unai Alvarez-Rodriguez, Federico Battiston, Guilherme Ferraz de Arruda, Yamir Moreno, Matjaž Perc, and Vito Latora. Evolutionary dynamics of higher-order interactions in social networks. *Nature Human Behaviour*, 5(5):586–595, 2021.

[7] Anthony B. Atkinson, Thomas Piketty, and Emmanuel Saez. Top incomes in the long run of history. *Journal of Economic Literature*, 49(1):3–71, March 2011. doi:10.1257/jel.49.1.3. URL `https://www.aeaweb.org/articles?id=10.1257/jel.49.1.3`.

[8] David H Autor. Skills, education, and the rise of earnings inequality among the "other 99 percent". *Science*, 344(6186):843–851, 2014.

[9] Jon Bakija, Adam Cole, and Bradley Heim. Jobs and income growth of top earners and the causes of changing income inequality: Evidence from u.s. tax return data. Mar 2012.

[10] Scott Barrett. Coordination vs. voluntarism and enforcement in sustaining international environmental cooperation. *Proceedings of the National Academy of Sciences*, 113(51):14515–14522, 2016.

[11] Andrew R Barron and Thomas M Cover. A bound on the financial value of information. *IEEE Transactions on Information Theory*, 34(5):1097–1100, 1988.

[12] Federico Battiston, Enrico Amico, Alain Barrat, Ginestra Bianconi, Guilherme Ferraz de Arruda, Benedetta Franceschiello, Iacopo Iacopini, Sonia Kéfi, Vito Latora, Yamir Moreno, et al. The physics of higher-order interactions in complex systems. *Nature Physics*, 17(10):1093–1098, 2021.

[13] Dominik Baumann, Erfaun Noorani, James Price, Ole Peters, Colm Connaughton, and Thomas B Schön. Non-ergodicity in reinforcement learning: robustness via ergodicity transformations. *arXiv preprint arXiv:2310.11335*, 2023.

[14] Luc Bauwens, Sébastien Laurent, and Jeroen VK Rombouts. Multivariate garch models: a survey. *Journal of applied econometrics*, 21(1):79–109, 2006.

[15] Hannah M Bayer and Paul W Glimcher. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron*, 47(1):129–141, 2005.

[16] Gary S Becker, Tomas J Philipson, and Rodrigo R Soares. The quantity and quality of life and the evolution of world inequality. *American economic review*, 95(1):277–291, 2005.

[17] Timothy EJ Behrens, Mark W Woolrich, Mark E Walton, and Matthew FS Rushworth. Learning the value of information in an uncertain world. *Nature neuroscience*, 10(9): 1214–1221, 2007.

[18] Philippe Belley and Lance Lochner. The changing role of family income and ability in determining educational achievement. *Journal of Human capital*, 1(1):37–89, 2007.

[19] Lorenzo Giovanni Bellu and Paolo Liberati. Policy impacts on inequality – simple inequality measures. *EASYPol, Analytical tools. Policy Support Service, Policy Assistance Division, FAO*, 2006.

[20] Michael Bengfort, Horst Malchow, and Frank M Hilker. The fokker–planck law of diffusion and pattern formation in heterogeneous environments. *Journal of mathematical biology*, 73(3):683–704, 2016.

[21] Yonatan Berman, Eshel Ben-Jacob, and Yoash Shapira. The dynamics of wealth inequality and the effect of income distribution. *PloS one*, 11:e0154196, 04 2016. doi:10.1371/journal.pone.0154196.

[22] Yonatan Berman, Ole Peters, and Alexander Adamou. Wealth inequality and the ergodic hypothesis: Evidence from the united states. *Claremont McKenna College Robert Day School of Economics & Finance Research Paper Series*, 2020.

[23] Luís Bettencourt. Urban growth and the emergent statistics of cities. *Science Advances*, 6:eaat8812, 08 2020. doi:10.1126/sciadv.aat8812.

[24] Luís M. A. Bettencourt. *Introduction to Urban Science Evidence and Theory of Cities as Complex Systems*. The MIT Press, August 2021. ISBN 9780262046008.

[25] Luís MA Bettencourt. The rules of information aggregation and emergence of collective intelligent behavior. *Topics in Cognitive Science*, 1(4):598–620, 2009.

[26] Luís MA Bettencourt. The origins of scaling in cities. *science*, 340(6139):1438–1441, 2013.

[27] Luís M.A. Bettencourt. Towards a statistical mechanics of cities. *Comptes Rendus Physique*, 20:308–318, 2019.

[28] Luís MA Bettencourt. Towards a statistical mechanics of cities. *Comptes Rendus Physique*, 20(4):308–318, 2019.

[29] Luis MA Bettencourt and Ruy M Ribeiro. Real time bayesian estimation of the epidemic potential of emerging infectious diseases. *PloS one*, 3(5):e2185, 2008.

[30] Luís MA Bettencourt, José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B West. Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the national academy of sciences*, 104(17):7301–7306, 2007.

[31] Luis MA Bettencourt, Greg J Stephens, Michael I Ham, and Guenter W Gross. Functional structure of cortical neuronal networks grown in vitro. *Physical Review E*, 75 (2):021915, 2007.

[32] Luís MA Bettencourt, Vadas Gintautas, and Michael I Ham. Identification of functional information subgraphs in complex networks. *Physical review letters*, 100(23):238701, 2008.

[33] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.

[34] Max R Blouin and Roberto Serrano. A decentralized market with common values uncertainty: Non-steady states. *The Review of Economic Studies*, 68(2):323–346, 2001.

[35] Lawrence Blume and David Easley. Heterogeneity, selection, and wealth dynamics. *Annu. Rev. Econ.*, 2(1):425–450, 2010.

[36] Bruce Boghosian. Kinetics of wealth and the pareto law. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 89:042804, 04 2014.

[37] Tim Bollerslev, Robert F Engle, and Jeffrey M Wooldridge. A capital asset pricing model with time-varying covariances. *Journal of political Economy*, 96(1):116–131, 1988.

[38] Jean-Philippe Bouchaud. Wealth condensation in a simple model of economy. *Physica A: Statistical Mechanics and its Applications*, 282:536–545, 2000.

[39] Jean-Philippe Bouchaud. On growth-optimal tax rates and the issue of wealth inequalities. *Journal of Statistical Mechanics: Theory and Experiment*, 2015, 08 2015. doi:10.1088/1742-5468/2015/11/P11011.

[40] Breno Braga, Signe-Mary McKernan, Caroline Ratcliffe, and Sandy Baum. Wealth inequality is a barrier to education and social mobility. *Urban Institute: Elevate the Debate. https://www. urban. org/research/publication/wealth-inequality-barrier-education-and-social-mobility*, 2017.

[41] Lucian Buşoniu, Robert Babuška, and Bart De Schutter. Multi-agent reinforcement learning: An overview. *Innovations in multi-agent systems and applications-1*, pages 183–221, 2010.

[42] Marco Cagetti and Mariacristina De Nardi. Wealth inequality: Data and models. *Macroeconomic dynamics*, 12(S2):285–313, 2008.

[43] John O Campbell. Universal darwinism as a process of bayesian inference. *Frontiers in Systems Neuroscience*, 10:49, 2016.

[44] Claudio Castellano, Matteo Marsili, and Alessandro Vespignani. Nonequilibrium phase transition in a model for social influence. *Physical Review Letters*, 85(16):3536, 2000.

[45] Tiago V de V Cavalcanti, Kamiar Mohaddes, and Mehdi Raissi. Growth, development and natural resources: New evidence using a heterogeneous panel analysis. *The Quarterly Review of Economics and Finance*, 51(4):305–318, 2011.

[46] Giulia Cencetti, Federico Battiston, Bruno Lepri, and Márton Karsai. Temporal properties of higher-order interactions in social networks. *Scientific reports*, 11(1):7028, 2021.

[47] Anirban Chakraborti and Bikas Chakrabarti. Statistical mechanics of money: How saving propensity affects its distribution. *The European Physical Journal B - Condensed Matter and Complex Systems*, 17, 09 2000. doi:10.1007/s100510070173.

[48] Arjun Chandrasekaran, Deshraj Yadav, Prithvijit Chattopadhyay, Viraj Prabhu, and Devi Parikh. It takes two to tango: Towards theory of ai's mind. *arXiv preprint arXiv:1704.00717*, 2017.

[49] Subrahmanyan Chandrasekhar. Stochastic problems in physics and astronomy. *Reviews of modern physics*, 15(1):1, 1943.

[50] Raj Chetty, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623, 2014.

[51] Theodor Cimpeanu, Francisco C Santos, Luís Moniz Pereira, Tom Lenaerts, and The Anh Han. Artificial intelligence development races in heterogeneous settings. *Scientific reports*, 12(1):1723, 2022.

[52] Simon Ciranka and Wouter van den Bos. Adolescent risk-taking in the context of exploration and social influence. *Developmental Review*, 61:100979, 2021.

[53] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[54] Richard C Connor. The benefits of mutualism: a conceptual framework. *Biological Reviews*, 70(3):427–457, 1995.

[55] Guy A Cooper and Stuart A West. Division of labour and the evolution of extreme specialization. *Nature ecology & evolution*, 2(7):1161–1167, 2018.

[56] Juan Carlos Córdoba and Geneviève Verdier. Lucas vs. lucas: On inequality and growth. *Economic Growth*, 2007.

[57] Vincent D Costa, Valery L Tran, Janita Turchi, and Bruno B Averbeck. Reversal learning and dopamine: a bayesian perspective. *Journal of Neuroscience*, 35(6):2407–2416, 2015.

[58] Thomas M Cover and Joy A Thomas. Information theory and statistics. *Elements of information theory*, 1(1):279–335, 1991.

[59] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, July 2006. ISBN 0471241954.

[60] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory 2nd Edition (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, July 2006. ISBN 0471241954.

[61] Richard T Cox. Probability, frequency and reasonable expectation. *American journal of physics*, 14(1):1–13, 1946.

[62] Flavio Cunha and James J Heckman. The evolution of inequality, heterogeneity and uncertainty in labor earnings in the u.s. economy. (13526), October 2007. doi:10.3386/w13526. URL `http://www.nber.org/papers/w13526`.

[63] Jonas Dalege, Denny Borsboom, Frenk van Harreveld, and Han LJ van der Maas. The attitudinal entropy (ae) framework as a general theory of individual attitudes. *Psychological Inquiry*, 29(4):175–193, 2018.

[64] Jonas Dalege, Mirta Galesic, and Henrik Olsson. Networks of beliefs: An integrative theory of individual-and social-level belief dynamics. 2023.

[65] Pranav Dandekar, Ashish Goel, and David T Lee. Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110 (15):5791–5796, 2013.

[66] Anna-Maria D'Cruz, Michael E Ragozzino, Matthew W Mosconi, Mani N Pavuluri, and John A Sweeney. Human reversal learning under conditions of certain versus uncertain outcomes. *Neuroimage*, 56(1):315–322, 2011.

[67] Pierre Degond, Jian Guo Liu, and Christian Ringhofer. Evolution of the distribution of wealth in an economic environment driven by local nash equilibria. *Journal of Statistical Physics*, 154(3):751–780, 2014. ISSN 0022-4715. doi:10.1007/s10955-013-0888-4.

[68] Tiziana Di Matteo, Tomaso Aste, and ST Hyde. Exchanges in complex networks: income and wealth distributions. In *The physics of complex systems (New advances and perspectives)*, pages 435–442. IOS Press, 2004.

[69] Kevin J Dooley. A complex adaptive systems model of organization change. *Nonlinear dynamics, psychology, and life sciences*, 1:69–97, 1997.

[70] Darrell Duffie and Gustavo Manso. Information percolation in large markets. *American Economic Review*, 97(2):203–209, 2007.

[71] Bertram Düring, Daniel Matthes, and Giuseppe Toscani. Kinetic equations modelling wealth redistribution: a comparison of approaches. *Physical Review E*, 78(5):056103, 2008.

[72] Mary Jane West Eberhard. The evolution of social behavior by kin selection. *The Quarterly Review of Biology*, 50(1):1–33, 1975.

[73] Maria K Eckstein, Sarah L Master, Ronald E Dahl, Linda Wilbrecht, and Anne GE Collins. Reinforcement learning and bayesian inference provide complementary models for the unique advantage of adolescents in stochastic reversal. *Developmental Cognitive Neuroscience*, 55:101106, 2022.

[74] Glen H Elder, Monica Kirkpatrick Johnson, and Robert Crosnoe. The emergence and development of life course theory. In *Handbook of the life course*, pages 3–19. Springer, 2003.

[75] Gary W Evans. The environment of childhood poverty. *American psychologist*, 59(2): 77, 2004.

[76] George W Evans and Seppo Honkapohja. Learning and macroeconomics. *Annu. Rev. Econ.*, 1(1):421–449, 2009.

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

[78] Lorenzo Fant, Onofrio Mazzarisi, Emanuele Panizon, and Jacopo Grilli. Stable cooperation emerges in stochastic multiplicative growth. *arXiv preprint arXiv:2202.02787*, 2022.

[79] Lorenzo Fant, Onofrio Mazzarisi, Emanuele Panizon, and Jacopo Grilli. Stable cooperation emerges in stochastic multiplicative growth. *Physical Review E*, 108(1):L012401, 2023.

[80] Roger Farmer and Jean-Philippe Bouchaud. Self-fulfilling prophecies, quasi non-ergodicity & wealth inequality. Technical report, National Bureau of Economic Research, 2020.

[81] Guilherme Ferraz de Arruda, Michele Tizzani, and Yamir Moreno. Phase transitions and stability of dynamical processes on hypergraphs. *Communications Physics*, 4(1): 24, 2021.

[82] Ronald Aylmer Fisher. *The genetical theory of natural selection.* , 1958.

[83] Jeffrey A Fletcher and Michael Doebeli. How altruism evolves: assortment and synergy. *Journal of evolutionary biology*, 19(5):1389–1393, 2006.

[84] Steven A Frank. Natural selection maximizes fisher information. *Journal of Evolutionary Biology*, 22(2):231–244, 2009.

[85] Steven A Frank. Natural selection. i. variable environments and uncertain returns on investment. *Journal of evolutionary biology*, 24(11):2299–2309, 2011.

[86] Steven A Frank. Natural selection. v. how to read the fundamental equations of evolutionary change in terms of information theory. *Journal of evolutionary biology*, 25 (12):2377–2396, 2012.

[87] Steven A Frank. Natural selection. iii. selection versus transmission and the levels of selection. *Journal of evolutionary biology*, 25(2):227–243, 2012.

[88] Steven A Frank. Natural selection. iv. the price equation. *Journal of evolutionary biology*, 25(6):1002–1019, 2012.

[89] Steven A Frank. Natural selection. v. how to read the fundamental equations of evolutionary change in terms of information theory. *Journal of evolutionary biology*, 25 (12):2377–2396, 2012.

[90] Willem E Frankenhuis, Karthik Panchanathan, and Daniel Nettle. Cognition in harsh and unpredictable environments. *Current Opinion in Psychology*, 7:76–80, 2016.

[91] Karl Friston. The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2):127–138, 2010.

[92] Drew Fudenberg and David K Levine. Whither game theory? towards a theory of learning in games. *Journal of Economic Perspectives*, 30(4):151–170, 2016.

[93] Drew Fudenberg and David K Levine. Whither game theory? towards a theory of learning in games. *Journal of Economic Perspectives*, 30(4):151–170, 2016.

[94] R Buckminster Fuller. *Synergetics: explorations in the geometry of thinking*. Estate of R. Buckminster Fuller, 1982.

[95] Xavier Gabaix, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. The dynamics of inequality. *Econometrica*, 84:2071–2111, 11 2016. doi:10.3982/ECTA13569.

[96] Mirta Galesic, Henrik Olsson, Jonas Dalege, Tamara Van Der Does, and Daniel L Stein. Integrating social and cognitive aspects of belief dynamics: towards a unifying framework. *Journal of the Royal Society Interface*, 18(176):20200857, 2021.

[97] Diego Garlaschelli and Maria I. Loffredo. Effects of network topology on wealth distributions. *Journal of Physics A*, 41:224018, 2008.

[98] Samuel J Gershman and Naoshige Uchida. Believing in dopamine. *Nature Reviews Neuroscience*, 20(11):703–714, 2019.

[99] Giovanni Giorgi and Chiara Gigliarano. The gini concentration index: A review of the inference literature. *Journal of Economic Surveys*, 31, 12 2016. doi:10.1111/joes.12185.

[100] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and economic behavior*, 22(1):1–29, 1998.

[101] Piotr J Gmytrasiewicz, Sanguk Noh, and Tad Kellogg. Bayesian update of recursive agent models. *User Modeling and User-Adapted Interaction*, 8:49–69, 1998.

[102] Claudia Goldin and Lawrence F Katz. The race between education and technology: The evolution of u.s. educational wage differentials, 1890 to 2005. (12984), March 2007. doi:10.3386/w12984.

[103] Jesús Gómez-Gardenes, Irene Reinares, Alex Arenas, and Luis Mario Floría. Evolution of cooperation in multiplex networks. *Scientific reports*, 2(1):1–6, 2012.

[104] Alison Gopnik and Henry M Wellman. Why the child's theory of mind really is a theory. 1992.

[105] Alison Gopnik, Shaun O'Grady, Christopher G Lucas, Thomas L Griffiths, Adrienne Wente, Sophie Bridgers, Rosie Aboody, Hoki Fung, and Ronald E Dahl. Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, 114(30): 7892–7899, 2017.

[106] Sanjeev Goyal and Fernando Vega-Redondo. Network formation and social coordination. *Games and Economic Behavior*, 50(2):178–207, 2005.

[107] Carlos Gracia-Lázaro, Alfredo Ferrer, Gonzalo Ruiz, Alfonso Tarancón, José A Cuesta, Angel Sánchez, and Yamir Moreno. Heterogeneous networks do not promote cooperation when humans play a prisoner's dilemma. *Proceedings of the National Academy of Sciences*, 109(32):12922–12926, 2012.

[108] Sanford Grossman. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of finance*, 31(2):573–585, 1976.

[109] Sanford J. Grossman and Joseph E. Stiglitz. On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3):393–408, 1980. ISSN 0002-8282. URL `https://www.jstor.org/stable/1805228`. Publisher: American Economic Association.

[110] Fatih Guvenen. Learning your earning: Are labor income shocks really very persistent? *American Economic Review*, 97 (3):687–712, 2007.

[111] Fatih Guvenen. Macroeconomics with heterogeneity: A practical guide. *National Bureau of Economic Research*, 2011.

[112] Daniel A Hackman, Martha J Farah, and Michael J Meaney. Socioeconomic status and the brain: mechanistic insights from human and animal research. *Nature reviews neuroscience*, 11(9):651–659, 2010.

[113] William D Hamilton. The evolution of altruistic behavior. *The American Naturalist*, 97(896):354–356, 1963.

[114] Peter Hammerstein and Reinhard Selten. Game theory and evolutionary biology. *Handbook of game theory with economic applications*, 2:929–993, 1994.

[115] The Anh Han, Tom Lenaerts, Francisco C Santos, and Luis Moniz Pereira. Voluntary safety commitments provide an escape from over-regulation in ai development. *Technology in Society*, 68:101843, 2022.

[116] Anthony Hannagan and Jonathan Morduch. Income gains and month-to-month income volatility: Household evidence from the us financial diaries. *NYU Wagner research paper*, (2659883), 2015.

[117] Eric A Hanushek and Ludger Woessmann. Education and economic growth. *Economics of education*, pages 60–67, 2010.

[118] Friedrich August Hayek. The use of knowledge in society. In *Modern Understandings of Liberty and Property*, pages 27–38. Routledge, 2013.

[119] Elisa Heinrich Mora, Cate Heine, Jacob J Jackson, Geoffrey B West, Vicky Chuqiao Yang, and Christopher P Kempes. Scaling of urban income inequality in the usa. *Journal of the Royal Society Interface*, 18(181):20210223, 2021.

[120] Els Heinsalu and Marco Patriarca. Kinetic models of immediate exchange. *The European Physical Journal B*, 87:1–10, 2014.

[121] Peter Helmberger and Sidney Hoos. Cooperative enterprise and organization theory. *Journal of farm economics*, 44(2):275–290, 1962.

[122] Ralph Hertwig, Greg Barron, Elke U Weber, and Ido Erev. Decisions from experience and the effect of rare events in risky choice. *Psychological science*, 15(8):534–539, 2004.

[123] Martin Hilbert. The more you know, the more you can grow: an information theoretic approach to growth in the information age. *Entropy*, 19(2):82, 2017.

[124] Yu Hsing. Economic growth and income inequality: the case of the us. *International Journal of Social Economics*, 32(7):639–647, 2005.

[125] Joachim Hubmer, Per Krusell, and Anthony A Smith Jr. Sources of us wealth inequality: Past, present, and future. *NBER Macroeconomics Annual*, 35(1):391–455, 2021.

[126] David Hume. *A treatise of human nature*. Courier Corporation, 2003.

[127] Tanya Ignatenko, Kirill Kondrashov, Marco Cox, and Bert de Vries. On preference learning based on sequential bayesian optimization with pairwise comparison. *arXiv preprint arXiv:2103.13192*, 2021.

[128] Alicia Izquierdo, Jonathan L Brigman, Anna K Radke, Peter H Rudebeck, and Andrew Holmes. The neural basis of reversal learning: an updated perspective. *Neuroscience*, 345:12–26, 2017.

[129] Matthew O Jackson and Yves Zenou. Games on networks. In *Handbook of game theory with economic applications*, volume 4, pages 95–163. Elsevier, 2015.

[130] Abel Jansma. Higher-order in-and-outeractions reveal synergy and logical dependence beyond shannon-information. *arXiv preprint arXiv:2205.04440*, 2022.

[131] Andrew H Jazwinski. *Stochastic processes and filtering theory*. Courier Corporation, 2007.

[132] Gregory A Johnson. Organizational structure and scalar stress. *Theory and explanation in archaeology*, pages 389–421, 1982.

[133] Laurent Valentin Jospin, Hamid Laga, Farid Boussaid, Wray Buntine, and Mohammed Bennamoun. Hands-on bayesian neural networks—a tutorial for deep learning users. *IEEE Computational Intelligence Magazine*, 17(2):29–48, 2022.

[134] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[135] Lisa A Keister. *Wealth in America: Trends in wealth inequality*. Cambridge University Press, 2000.

[136] J. L. Kelly. A new interpretation of information rate. *IRE Trans. Inf. Theory*, 2: 185–189, 1956.

[137] John L Kelly. A new interpretation of information rate. *the bell system technical journal*, 35(4):917–926, 1956.

[138] Jordan Kemp and Luis Bettencourt. Bayesian origins of growth, cooperation, and inequality in populations of learning agents. *Bulletin of the American Physical Society*, 2023.

[139] Jordan Kemp, Max-Olivier Hongler, and Olivier Gallay. Stochastic pairwise preference convergence in bayesian agents. *Physical Review E*, 109(5):054106, 2024.

[140] Jordan T Kemp and Luís MA Bettencourt. Statistical dynamics of wealth inequality in stochastic models of growth. *Physica A: Statistical Mechanics and its Applications*, 607:128180, 2022.

[141] Jordan T Kemp and Luís MA Bettencourt. Learning increases growth and reduces inequality in shared noisy environments. *PNAS Nexus*, page pgad093, 2023.

[142] Jordan T Kemp, Adam G Kline, and Luís Bettencourt. Information synergy maximizes the growth rate of heterogeneous groups. *arXiv preprint arXiv:2307.01380*, 2023.

[143] Jordan T Kemp, Adam G Kline, and Luís MA Bettencourt. Information synergy maximizes the growth rate of heterogeneous groups. *PNAS nexus*, page pgae072, 2024.

[144] Mohammad Emtiyaz Khan, Young Jun Ko, and Matthias Seeger. Scalable collaborative bayesian preference learning. In *Artificial Intelligence and Statistics*, pages 475–483. PMLR, 2014.

[145] Celeste Kidd and Benjamin Y Hayden. The psychology and neuroscience of curiosity. *Neuron*, 88(3):449–460, 2015.

[146] Shigeo S Kimura, Kohta Murase, Imre Bartos, Kunihito Ioka, Ik Siong Heng, and Peter Mészáros. Transejecta high-energy neutrino emission from binary neutron star mergers. *Physical Review D*, 98(4):043020, 2018.

[147] Kevin B Korb and Ann E Nicholson. *Bayesian artificial intelligence*. CRC press, 2010.

[148] Alan B Krueger and Mikael Lindahl. Education for growth: Why and for whom? *Journal of economic literature*, 39(4):1101–1136, 2001.

[149] Aanjaneya Kumar, Valerio Capraro, and Matjaž Perc. The evolution of trust and trustworthiness. *Journal of the Royal Society Interface*, 17(169):20200491, 2020.

[150] Aanjaneya Kumar, Valerio Capraro, and Matjaž Perc. The evolution of trust and trustworthiness. *Journal of the Royal Society Interface*, 17(169):20200491, 2020.

[151] Mordecai Kurz and Maurizio Motolese. Endogenous uncertainty and market volatility. *Economic Theory*, 17:497–544, 2001.

[152] Edo Kussell and Stanislas Leibler. Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309(5743):2075–2078, 2005.

[153] Paul R Lawrence and Jay W Lorsch. Differentiation and integration in complex organizations. *Administrative science quarterly*, pages 1–47, 1967.

[154] Stephen Le and Robert Boyd. Evolutionary dynamics of the continuous iterated prisoner's dilemma. *Journal of theoretical biology*, 245(2):258–267, 2007.

[155] Kyu-Min Lee, Jae-Suk Yang, Gunn Kim, Jaesung Lee, Kwang-Il Goh, and In-mook Kim. Impact of the topology of global macroeconomic network on the spreading of economic crises. *PloS one*, 6(3):e18443, 2011.

[156] Michael D Lee and Eric-Jan Wagenmakers. *Bayesian cognitive modeling: A practical course*. Cambridge university press, 2014.

[157] Daniel A Levinthal. Adaptation on rugged landscapes. *Management science*, 43(7): 934–950, 1997.

[158] Jie Li and Bruce M. Boghosian. Duality in an asset exchange model for wealth distribution. *Physica A: Statistical Mechanics and its Applications*, 497:154–165, 2018. doi:10.1016/j.physa.2018.10.042.

[159] Jie Li, Bruce Boghosian, and Chengli Li. The affine wealth model: An agent-based model of asset exchange that allows for negative-wealth agents and its empirical validation. *Physica A: Statistical Mechanics and its Applications*, 516, 10 2018. doi:10.1016/j.physa.2018.10.042.

[160] Aaron D Lightner, Anne C Pisor, and Edward H Hagen. In need-based sharing, sharing is more important than need. *Evolution and Human Behavior*, 2023.

[161] Meng Liu and Ke Wang. Stochastic lotka–volterra systems with lévy noise. *Journal of Mathematical Analysis and Applications*, 410(2):750–763, 2014.

[162] Charles G Lord, Lee Ross, and Mark R Lepper. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology*, 37(11):2098, 1979.

[163] Michael F Lovenheim. The effect of liquid housing wealth on college enrollment. *Journal of Labor Economics*, 29(4):741–771, 2011.

[164] Robert E. Lucas. The industrial revolution: past and future. *2003 Annual Report of the Federal Reserve Bank of Minneapolis*, pages 5,20, 2004.

[165] Artur Luczak, Bruce L McNaughton, and Yoshimasa Kubo. Neurons learn by predicting future activity. *Nature machine intelligence*, 4(1):62–72, 2022.

[166] David JC MacKay. Bayesian neural networks and density networks. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 354(1):73–80, 1995.

[167] James G March. Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87, 1991.

[168] Kevin A McCabe, Stephen J Rassenti, and Vernon L Smith. Game theory and reciprocity in some extensive form experimental games. *Proceedings of the National Academy of Sciences*, 93(23):13421–13428, 1996.

[169] Pedro AM Mediano, Fernando E Rosas, Andrea I Luppi, Henrik J Jensen, Anil K Seth, Adam B Barrett, Robin L Carhart-Harris, and Daniel Bor. Greater than the parts: a review of the information decomposition approach to causal emergence. *Philosophical Transactions of the Royal Society A*, 380(2227):20210246, 2022.

[170] Matus Medo. Breakdown of the mean-field approximation in a wealth distribution model. *arXiv.org, Quantitative Finance Papers*, 2009, 09 2008. doi:10.1088/1742-5468/2009/02/P02014.

[171] Costas Meghir and Luigi Pistaferri. Earnings, consumption and life cycle choices. In *Handbook of labor economics*, volume 4, pages 773–854. Elsevier, 2011.

[172] Fabio Milani. Expectations, learning and macroeconomic persistence. *Journal of monetary Economics*, 54(7):2065–2082, 2007.

[173] Paul R Milgrom. Rational expectations, information acquisition, and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 921–943, 1981.

[174] George A Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*, 63(2):81, 1956.

[175] T. M. Mitchell. *Machine Learning*. McGraw-Hill Education, March 1997. ISBN 0070428077.

[176] Paul Morris. Asia's four little tigers: a comparison of the role of education in their development. *Comparative education*, 32(1):95–110, 1996.

[177] mpiktas (https://stats.stackexchange.com/users/2116 /mpiktas). Variance of a function of one random variable. Cross Validated. URL `https://stats.stackexchange.com/q/5790`. URL:https://stats.stackexchange.com/q/5790 (version: 2020-03-01).

[178] Kevin P Murphy. Conjugate bayesian analysis of the gaussian distribution. *def*, 1 ($2\sigma 2$):16, 2007.

[179] Daniel B Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the econometric society*, pages 347–370, 1991.

[180] Martin A Nowak. Five rules for the evolution of cooperation. *science*, 314(5805): 1560–1563, 2006.

[181] Martin A Nowak, Alex McAvoy, Benjamin Allen, and Edward O Wilson. The general form of hamilton's rule makes no predictions and cannot be tested empirically. *Proceedings of the National Academy of Sciences*, 114(22):5665–5670, 2017.

[182] Ndidi Bianca Ogbo, Aiman Elragig, and The Anh Han. Evolution of coordination in pairwise and multi-player interactions via prior commitments. *Adaptive Behavior*, 30 (3):257–277, 2022.

[183] Hisashi Ohtsuki, Christoph Hauert, Erez Lieberman, and Martin A Nowak. A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092): 502–505, 2006.

[184] Eckehard Olbrich, Nils Bertschinger, and Johannes Rauh. Information decomposition and synergy. *Entropy*, 17(5):3501–3517, 2015.

[185] Zita Oravecz, Francis Tuerlinckx, and Joachim Vandekerckhove. Bayesian data analysis with the bivariate hierarchical ornstein-uhlenbeck process model. *Multivariate behavioral research*, 51(1):106–119, 2016.

[186] Scott E. Page. *The difference: how the power of diversity creates better groups, firms, schools, and societies*. Princeton Univ. Press, Princeton, NJ, 3. print., and 1. paperback print., with a new preface edition, 2007. ISBN 978-0-691-13854-1 978-0-691-12838-2.

[187] Josefine Pallavicini, Bjørn Hallsson, and Klemens Kappel. Polarization in groups of bayesian agents. *Synthese*, 198:1–55, 2021.

[188] M. Patriarca, E. Heinsalu, and A. Chakraborti. Basic kinetic wealth-exchange models: common features and open problems. *The European Physical Journal B*, 73:145–153, 2006.

[189] Michael B Paulsen. The economics of human capital and investment in higher education. *The finance of higher education: Theory, research, policy, and practice*, pages 55–94, 2001.

[190] Elizabeth Pennisi. On the origin of cooperation, 2009.

[191] John W Pepper. Relatedness in trait group models of social evolution. *Journal of Theoretical Biology*, 206(3):355–368, 2000.

[192] John W Pepper and Barbara B Smuts. A mechanism for the evolution of altruism among nonkin: positive assortment through environmental feedback. *The American Naturalist*, 160(2):205–213, 2002.

[193] Matjaž Perc, Jesús Gómez-Gardenes, Attila Szolnoki, Luis M Floría, and Yamir Moreno. Evolutionary dynamics of group interactions on structured populations: a review. *Journal of the royal society interface*, 10(80):20120997, 2013.

[194] Matjaž Perc, Jillian J Jordan, David G Rand, Zhen Wang, Stefano Boccaletti, and Attila Szolnoki. Statistical physics of human cooperation. *Physics Reports*, 687:1–51, 2017.

[195] Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15:1216–1221, 12 2019. doi:10.1038/s41567-019-0732-0.

[196] Ole Peters and Alexander Adamou. The ergodicity solution of the cooperation puzzle. *Philosophical Transactions of the Royal Society A*, 380(2227):20200425, 2022.

[197] Thomas Piketty and Emmanuel Saez. Inequality in the long run. *Science*, 343(6186): 838–843, 2014.

[198] George R Price. Extension of covariance selection mathematics. *Annals of human genetics*, 35(4):485–490, 1972.

[199] George R Price. Fisher's 'fundamental theorem'made clear. *Annals of human genetics*, 36(2):129–140, 1972.

[200] George R Price et al. Selection and covariance. *Nature*, 227:520–521, 1970.

[201] David C Queller. Kinship, reciprocity and synergism in the evolution of social behaviour. *Nature*, 318(6044):366–367, 1985.

[202] Roberta Raileanu, Emily Denton, Arthur Szlam, and Rob Fergus. Modeling others using oneself in multi-agent reinforcement learning. In *International conference on machine learning*, pages 4257–4266. PMLR, 2018.

[203] David G Rand, Samuel Arbesman, and Nicholas A Christakis. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences*, 108(48):19193–19198, 2011.

[204] Olivier Rivoire and Stanislas Leibler. The value of information for populations in varying environments. *Journal of Statistical Physics*, 142:1124–1166, 2011.

[205] Herbert Robbins. The empirical bayes approach to testing statistical hypotheses. *Revue de l'Institut International de Statistique*, pages 195–208, 1963.

[206] Gareth O Roberts, Omiros Papaspiliopoulos, and Petros Dellaportas. Bayesian inference for non-gaussian ornstein–uhlenbeck stochastic volatility processes. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(2):369–393, 2004.

[207] Joel L Sachs and James J Bull. Experimental evolution of conflict mediation between genomes. *Proceedings of the National Academy of Sciences*, 102(2):390–395, 2005.

[208] Joel L Sachs, Ulrich G Mueller, Thomas P Wilcox, and James J Bull. The evolution of cooperation. *The Quarterly review of biology*, 79(2):135–160, 2004.

[209] Anand Sahasranaman and Henrik Jeldtoft Jensen. Dynamics of reallocation within india's income distribution. *Indian Economic Review*, 56(1):1–23, 2021.

[210] A. I. Saichev, Y. Malevergne, and D. Sornette. Theory of zipf's law and beyond. *Springer Science Business Media*, 632, 2009.

[211] Alexander I. Saichev, Yannick Malevergne, and Didier Sornette. Theory of zipf's law and of general power law distributions with gibrat's law of proportional growth. *arXiv: General Finance*, 2008.

[212] Francisco C Santos, Jorge M Pacheco, and Brian Skyrms. Co-evolution of pre-play signaling and cooperation. *Journal of Theoretical Biology*, 274(1):30–35, 2011.

[213] Marten Scheffer, Bas Van Bavel, Ingrid Leemput, and Egbert Nes. Inequality in nature and society. *Proceedings of the National Academy of Sciences*, 114:201706412, 11 2017. doi:10.1073/pnas.1706412114.

[214] Laura Schenker. Performance games: Simulating the effect of incentives, monitoring, and agent adaptation on institutional data manipulation. 02 2023.

[215] Laura Schenker. Performance management in local bureaucracies: routines, replication, and resistance. 2023.

[216] Markus Schläpfer, Luís MA Bettencourt, Sébastian Grauwin, Mathias Raschke, Rob Claxton, Zbigniew Smoreda, Geoffrey B West, and Carlo Ratti. The scaling of human interactions with city size. *Journal of the Royal Society Interface*, 11(98):20130789, 2014.

[217] Elad Schneidman, William Bialek, and Michael J Berry. Synergy, redundancy, and independence in population codes. *Journal of Neuroscience*, 23(37):11539–11553, 2003.

[218] Theodore W Schultz. Investment in human capital. the role of education and of research. *ERIC*, 1971.

[219] Rüdiger J Seitz and Hans-Ferdinand Angel. Belief formation–a driving force for brain evolution. *Brain and Cognition*, 140:105548, 2020.

[220] Paulo Shakarian, Patrick Roos, and Anthony Johnson. A review of evolutionary graph theory with applications to game theory. *Biosystems*, 107(2):66–80, 2012.

[221] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. doi:10.1002/j.1538-7305.1948.tb01338.x.

[222] Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.

[223] Rajesh Singh, Dipanjan Ghosh, and R Adhikari. Fast bayesian inference of the multivariate ornstein-uhlenbeck process. *Physical Review E*, 98(1):012136, 2018.

[224] Lennart Sjögren. Brownian motion: Langevin equation, 2023. URL `http://web.phys.ntnu.no/~ingves/Teaching/TFY4275/Downloads/kap6.pdf`.

[225] Wataru Souma, Yoshi Fujiwara, and Hideaki Aoyama. Small-world effects in wealth distribution. *arXiv preprint cond-mat/0108482*, 2001.

[226] B Spagnolo, A Fiasconaro, and D Valenti. Noise induced phenomena in lotka-volterra systems. *Fluctuation and Noise Letters*, 3(02):L177–L185, 2003.

[227] Vivek H Sridhar, Liang Li, Dan Gorbonos, Máté Nagy, Bianca R Schell, Timothy Sorochkin, Nir S Gov, and Iain D Couzin. The geometry of decision-making in individuals and collectives. *Proceedings of the National Academy of Sciences*, 118(50): e2102157118, 2021.

[228] Arthur L Stinchcombe. *Information and organizations*, volume 19. Univ of California Press, 1990.

[229] Arthur L Stinchcombe. Social structure and organizations. In *Economics meets sociology in strategic management*, volume 17, pages 229–259. Emerald Group Publishing Limited, 2000.

[230] Viktor Stojkoski, Trifce Sandev, Lasko Basnarkov, Ljupco Kocarev, and Ralf Metzler. Generalised geometric brownian motion: Theory and applications to option pricing. 10 2020.

[231] Viktor Stojkoski, Petar Jolakoski, Arnab Pal, Trifce Sandev, Ljupco Kocarev, and Ralf Metzler. Income inequality and mobility in geometric brownian motion with stochastic resetting: theoretical results and empirical evidence of non-ergodicity. *Philosophical Transactions of the Royal Society A*, 380(2224):20210157, 2022.

[232] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[233] John Sweller. Cognitive load during problem solving: Effects on learning. *Cognitive science*, 12(2):257–285, 1988.

[234] György Szabó and Gabor Fath. Evolutionary games on graphs. *Physics reports*, 446 (4-6):97–216, 2007.

[235] Misako Takayasu, Hayafumi Watanabe, and Hideki Takayasu. Generalised central limit theorems for growth rate distribution of complex systems. *Journal of Statistical Physics*, 155:47–71, 2013.

[236] Sebastian Thrun. Exploration in active learning. *Handbook of Brain Science and Neural Networks*, pages 381–384, 1995.

[237] Jarle Tufto. Covariance between a normally distributed variable and its exponent. Cross Validated. URL https://stats.stackexchange.com/q/274853. URL:https://stats.stackexchange.com/q/274853 (version: 2017-04-20).

[238] Thomas F Varley and Erik Hoel. Emergence as the conversion of information: A unifying theory. *Philosophical Transactions of the Royal Society A*, 380(2227):20210150, 2022.

[239] Thomas F Varley, Maria Pope, Joshua Faskowitz, and Olaf Sporns. Multivariate information theory uncovers synergistic subsystems of the human cerebral cortex. *Communications Biology*, 6(1):451, 2023.

[240] Iris Vilares and Konrad Kording. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, 1224 (1):22–39, 2011.

[241] Friederike Wall. Agent-based modeling in managerial science: an illustrative survey and study. *Review of Managerial Science*, 10(1):135–193, 2016.

[242] Gian-Reto Walther. Community and ecosystem responses to recent climate change. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1549):2019–2024, 2010.

[243] David G Weissman, Mark Hatzenbuehler, Mina Cikara, Deanna Barch, and PhD McLaughlin, Katie A. Antipoverty programs mitigate socioeconomic disparities in brain structure and psychopathology among u.s. youths. *PsyArXiv*, Nov 2021. doi:10.31234/osf.io/8nhej. URL psyarxiv.com/8nhej.

[244] Robert Wilson. Incentive efficiency of double auctions. *Econometrica: Journal of the Econometric Society*, pages 1101–1115, 1985.

[245] Justin Wolfers and Eric Zitzewitz. Prediction Markets. *Journal of Economic Perspectives*, 18(2):107–126, June 2004. ISSN 0895-3309. doi:10.1257/0895330041371321. URL https://www.aeaweb.org/articles?id=10.1257/0895330041371321.

[246] Asher Wolinsky. Information revelation in a market with pairwise meetings. *Econometrica: Journal of the Econometric Society*, pages 1–23, 1990.

[247] Charley M Wu, Eric Schulz, Maarten Speekenbrink, Jonathan D Nelson, and Björn Meder. Generalization guides human exploration in vast decision spaces. *Nature human behaviour*, 2(12):915–924, 2018.

[248] Sarah A Wu, Rose E Wang, James A Evans, Joshua B Tenenbaum, David C Parkes, and Max Kleiman-Weiner. Too many cooks: Bayesian inference for coordinating multi-agent collaboration. *Topics in Cognitive Science*, 13(2):414–432, 2021.

[249] Matthieu Wyart and Jean-Philippe Bouchaud. Statistical models for company growth. *Physica A-statistical Mechanics and Its Applications*, 326:241–255, 2003.