



UvA-DARE (Digital Academic Repository)

Emergent poverty traps at multiple levels impede social mobility

Dupont, C.; Roy, D.

DOI

[10.1057/s41599-025-06089-9](https://doi.org/10.1057/s41599-025-06089-9)

Publication date

2025

Document Version

Final published version

Published in

Humanities and Social Sciences Communications

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Dupont, C., & Roy, D. (2025). Emergent poverty traps at multiple levels impede social mobility. *Humanities and Social Sciences Communications*, 12, Article 1777. <https://doi.org/10.1057/s41599-025-06089-9>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.



ARTICLE



<https://doi.org/10.1057/s41599-025-06089-9>

OPEN

Emergent poverty traps at multiple levels impede social mobility

Charles Dupont¹ & Debraj Roy²✉

Eradicating extreme poverty and inequality are the key leverage points to achieve the seventeen Sustainable Development Goals (SDGs). However, the reduction of extreme poverty and inequality remains vulnerable to shocks such as pandemics and climate change. Numerous models have been developed to examine the complex social interactions giving rise to inequality and persistent poverty, yet few approaches include multilevel dynamics. Here, we introduce a heterogeneous agent-based model to identify conditions underlying poverty traps at different levels, which manifest as distinct statistical steady-state outcomes. We find that vulnerabilities emerge from the interaction between individual and institutional mechanisms. Individual characteristics like risk aversion, attention, and saving propensity can lead to sub-optimal diversification and low capital accumulation. These individual drivers are reinforced by institutional mechanisms such as lack of financial inclusion, access to technology, and economic segregation, leading to persistent inequality and poverty traps. Our experiments demonstrate that addressing the above factors yields a “double dividend”—reducing poverty and inequality within and between communities and creating positive feedback that can withstand shocks. Finally, we demonstrate that our theoretical model can be used as a sandbox for cost-benefit analysis of intervention strategies.

¹University of Amsterdam, Amsterdam, Netherlands. ²Univeristy of Amsterdam, Amsterdam, Netherlands. ✉email: d.roy@uva.nl

Introduction

At the forefront of the Sustainable Development Goals lies one of the most important and unresolved global challenges of our time: eradicating extreme poverty and reducing inequality within and among countries. Although poverty rates have decreased over the years, nearly 10% of the world's population still lives in extreme poverty, characterized by a lack of basic needs such as access to food, clean drinking water, education, and information (Hasell et al., 2022). Emerging yet inconclusive evidence suggests that COVID-19 may have been a factor in this positive trend (Chancel et al., 2021; Yonzan et al., 2023). Estimates for 2023 indicate that poverty rates have likely returned to 2019 levels, with about 691 million people (8.6% of the world population) living in extreme poverty (Yonzan et al., 2023). The pandemic has also caused the largest increase in inequality between countries in three decades, and reducing inequality within and between countries is a scientific and a policy challenge (Chancel et al., 2021). In addition, climate change is proving to be a major obstacle for impoverished households and communities to adapt and develop in a sustainable way (Chancel et al., 2021). A recent review by Méjean et al. has indicated that the poor are expected to be disproportionately affected by the increasing frequency and intensity of climate-related shocks, potentially reversing progress towards poverty reduction and further increasing economic inequality (Méjean et al., 2024).

High levels of economic inequality can exacerbate poverty by concentrating wealth and opportunities among a small segment of the population, leaving others with limited resources and prospects for advancement. Poverty can act as a barrier to economic mobility—understood here as the ability of individuals and communities to improve their economic standing and escape low-income circumstances over time, rather than strictly as rank-order changes within a fixed distribution—as individuals born into poverty often face systemic obstacles such as inadequate access to education, healthcare, and employment opportunities. Limited economic mobility can perpetuate cycles of poverty, where individuals and possibly entire communities remain trapped in low-income circumstances across generations. Reflecting on the trajectory of economic development reveals a transition from a state characterized by extreme poverty to one marked by pervasive inequality. The empirical evidence highlights two key facts.

Multi-level poverty traps. The alleviation of extreme poverty has progressed at a slow pace in Sub-Saharan Africa, and trends indicate an increase in poverty levels within the regions of the Middle East and North Africa, suggesting the presence of “poverty traps” (see Appendix A for a brief review). This trap involves self-reinforcing mechanisms that keep regions, countries, and neighbourhoods in a cycle of poverty (Azariadis and Stachurski, 2005). Barrett et al. (Barrett and Carter, 2013; Barrett and Swallow, 2006; Barrett et al., 2016; Santos and Barrett, 2018) and Adato et al. (Adato et al., 2013) provide evidence of multiple equilibrium in asset dynamics in rural Kenya, Madagascar, and South Africa. Evidence from rural China (Jalan and Ravallion, 2002) and Bangladesh (Balboni et al., 2022) shows that living in poor areas diminishes the productivity of a farmer's investments, with inadequate community capital, such as limited infrastructure, potentially trapping households in poverty. An empirical survey (Kraay and McKenzie, 2014) examined mechanisms like S-shaped savings functions, coordination failures, hunger-based traps, and occupational poverty traps. The survey suggests that such traps are rare and mostly confined to remote or disadvantaged areas (Kraay and McKenzie, 2014). Emerging research highlights behavioural and geographic poverty

traps as more plausible (Kraay and McKenzie, 2014). Recent research has observed that poverty traps often operate across multiple levels, prompting the conceptualization of multi-level or fractal poverty traps (Barrett and Swallow, 2006). The implication of the poverty trap at multiple levels for policy is significant, suggesting that cross-level interactions can reinforce or mitigate poverty traps (Barrett et al., 2018; Radosavljevic et al., 2021).

Persistent horizontal inequality. Inequality appears at various spatial scales, from global to local, affecting regions, countries, cities, and neighbourhoods differently. In recent decades, while global inequalities between countries have decreased, inequalities within countries have surged. For example, China's sudden shift from extreme poverty to extreme inequality is far from being an exception, as many other countries today face similar challenges (World Bank and the Development Research Center of the State Council, The People's Republic of China, 2022; Qureshi, 2023; Xie and Zhou, 2014). Within cities, inequality can be stark between different neighbourhoods (e.g., formal vs. informal (Mutlu et al., 2023, Roy et al., 2018b)) and group identities (e.g., ethnicity, religion (Roy et al., 2018a)). There are many persistent horizontal inequalities in developing countries, such as northern groups in Nigeria or Ghana, Somalis in Kenya, Hutus in Rwanda and Burundi, or Muslims in India. Inequality can stem from differences in education (Galor and Zeira, 1993, Stiglitz, 1973), access to resources (Goderis and Malone, 2011) and technology (Aghion et al., 1999), individual risk aversion (Aghion and Bolton, 1992, Banerjee and Newman, 1991, Mengesha and Roy, 2025) and saving propensity (Banerjee and Yakovenko, 2010, Chakraborti and Chakrabarti, 2000), globalization and government policies (Chong and Gradstein, 2007). In conclusion, poverty traps and inequality are deeply intertwined, with high inequality often contributing to the persistence of poverty traps.

Although there is clear evidence of poverty traps and persistent inequality at multiple levels, the impact of cross-level interactions in poverty and emerging inequality is not well understood. The concept of multiple equilibria in current poverty traps literature typically describes scenarios where different initial conditions lead to divergent outcomes. However, this contrasts with interpretations in multi-agent game-theoretic models, where identical initial conditions can produce multiple distinct outcomes due to factors such as strategic complementarities, coordination failures, and macroeconomic fluctuations. Furthermore, current models attribute income variation primarily to differences in preferences and endowments, neglecting the role of variations in risky outcomes. This approach is logically incomplete, as it fails to explore what the long-term wealth distribution would look like in a scenario where agents start with identical endowments and preferences.

While a number of agent-based models have examined inequality and poverty, our understanding of the co-evolution of poverty traps across individual and community levels, driven by behavioural heterogeneity, remains a key challenge. Fierro et al. have investigated how monetary policy is impacted by income inequality using a macroscopic agent-based model (Fierro et al., 2023). Caiani et al. analysed the relationship between inequality and economic development with an agent-based stock-flow consistent model with different worker classes (Caiani et al., 2019). Botta et al. examined interactions between finance, instability, and inequality, also using an agent-based stock-flow consistent macroeconomic model (Botta et al., 2021). Palagi et al. studied interactions between inequality, income dynamics, and credit constraints in a generational agent-based model of households (Palagi et al., 2023). Although valuable in illuminating

different facets of poverty and inequality, these approaches do not explicitly model emergent multi-level dynamics of poverty traps arising from interactions across different scales.

In response to these challenges, we have developed a two-level agent-based model that is novel in its study of the linkages between heterogeneous individual behaviour, community structure, and differential institutional access on emergent inequality and poverty at both individual and community levels. Individual agents are connected in a social network from which we extract communities. We construct an economy where heterogeneous agents experience uninsurable idiosyncratic shocks and smooth consumption. Agents are risk-averse and have identical preferences deriving utility from consumption and can engage in two types of productive projects: one safe but low-return project and several high-risk, high-return projects. Agents within the same community can form self-financing groups and undertake joint risk-sharing investment projects. Agents have an attention mechanism that allows agents to periodically update their investment portfolios according to observed realisations of project returns (see “Methods” section for details). In order to study the model’s behaviour, we construct Monte Carlo (Sobol) samples and run several repetitions for each combination of parameters in order to capture stochastic effects (Saltelli, 2002) (see “Experiments” in “Methods” for details). Then, we categorise the model runs at both the micro (agent) and meso (community) levels according to the evolution of wealth trajectories, which tend towards characteristic long-run steady states for the given parameter sets. We find robust evidence with all experimental setups for the existence of three distinct regimes: a persistent poverty trap for all (“All Poor”), a state of high inequality where some escape poverty (“Some Rich”), and a rarer scenario where all prosper (“All Rich”). The application of a manifold learning-based GSA method (see Appendix B) allows us to identify critical parameter regions responsible for giving rise to each regime, and we observe important differences in the sensitivity indices of these parameters at different temporal scales as well as micro (agent), meso (community), and macro (population) scales. Intervention experiments indicate that it is not possible to extricate agents from poverty in the All Poor regime, while interventions of sufficient magnitude can help to alleviate poverty to some degree for the Some Rich regime. We find key individual (behavioural) and community (institutional) mechanisms that are important to achieve sustainable reduction in poverty and inequality. At the individual level, behavioural characteristics like risk aversion, attention, and saving propensity can lead to sub-optimal diversification and low capital accumulation. At the community level, institutional drivers such as lack of financial inclusion, access to technology, and economic segregation are key drivers of inequality and poverty traps. Our results show that targeting horizontal inequality (between-community inequality) increases the rate of successful interventions through feedback.

Methods

In this section, we describe the key components of the agent-based model as well as relevant theory.

Agent social network. Homophily, the tendency for similar agents to form social connections, has been shown to play a significant role in how social networks are organised. Indeed, most people’s social networks are relatively homogeneous with respect to socioeconomic and demographic characteristics such as race and ethnicity, age, education, and household income (McPherson et al., 2001). This has significant implications for the kind of information that people have access to, as well as the behaviours that we tend to adopt. In a recent paper, Talaga and

Nowak propose a Social Distance Attachment (SDA) model for randomly constructing social networks based on homophily (Talaga and Nowak, 2020). Furthermore, they show that their model is able to reproduce many of the characteristic features of social networks, making it an appealing choice for our present purposes. Existing poverty trap literature (Chantarat and Barrett, 2012) highlights the complex relationship between social networks, wealth distribution, and poverty alleviation, emphasizing both the potential benefits and limitations of social network capital in addressing persistent poverty.

Here, we construct an undirected social network consisting of $N = 1225$ agents, such that each agent possesses some initial amount of wealth drawn from a normal distribution centred at $\mu = 10$ with standard deviation $\sigma = 1$. An edge is added between two agents with respective wealth levels w_i and w_j with probability

$$p_{ij} = \frac{1}{1 + [b^{-1}d(w_i, w_j)]^\alpha}, \quad (1)$$

where $d(w_i, w_j) = |w_i - w_j|$, b is the characteristic distance at which $p_{ij} = 0.5$, and α is the homophily parameter (a higher value means that attachment is more likely between agents with similar wealth levels) (Talaga and Nowak, 2020). In practice we set the characteristic distance, based on preliminary results, to the average pairwise (wealth) distance between all agents divided by 15.

Due to the probabilistic nature of edge creation, this algorithm may yield multiple disconnected components. If so, we ensure connectedness as follows: (1) identify the largest connected component and (2) for every other connected component, pick a random node and insert an edge between this node and whichever node inside the largest connected component is closest in wealth. Figure 1a provides an example of a small SDA graph including agent wealth levels to highlight the way in which edges are formed based on the property of homophily. The reader may refer to Fig. 3 of Appendix C for some key distributions (e.g., number of communities, community sizes) for the SDA graphs that were generated with $N = 1225$ agents.

Agent community. The formation of communities based on homophily is one of the most important characteristics of social networks. In our model, a community is defined as a subset of nodes in a graph such that “connections between those nodes are denser than connections with the rest of the network”, and community structure significantly impacts the types of interactions and dynamics that can occur in a population (Radicchi et al., 2004). Concretely, agents within the same community typically interact and share information more often than with agents in different communities, which has important consequences for how a system evolves over time. Extracting a network’s community structure is therefore an important problem, which we tackle by using a semi-synchronous label propagation algorithm proposed by Cordasco and Gargano (Cordasco and Gargano, 2011). This approach was chosen as it relies on how information diffuses throughout a network to identify communities, is efficient, and shown to always converge to a stable labelling (Cordasco and Gargano, 2011).

Running the algorithm after graph construction yields a set C of communities, where each element is a set of agents and such that one agent cannot be in more than one community. However, we expand the notion of community membership as follows: an agent is a member of its own community as well as any community that agents it is adjacent to are part of. Figure 1b displays a stylised example of this procedure. Node 1 belongs to the blue community and the orange community since one of its

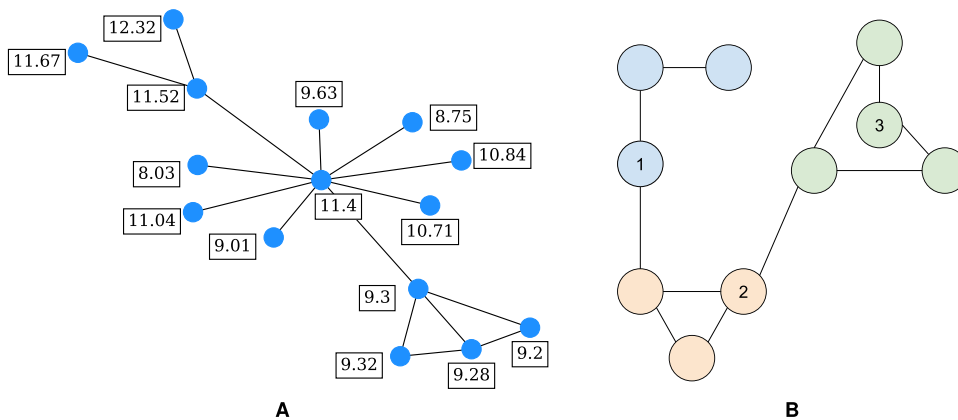


Fig. 1 Creation of social network and community. **A** SDA graph with homophily parameter $\alpha = 8$ for 15 agents with wealth values drawn from $\mathcal{N}(\mu = 10, \sigma = 1)$. **B** Toy example of community detection via label propagation. Colors correspond to the result of applying the label propagation algorithm. However, as described in this section, we consider an agent to be a member of its own community as well as a member of any community that agents that it is adjacent to are part of. Hence, node 1 belongs to the blue and orange communities, node 2 belongs to the orange and green communities, and node 3 belongs only to the green community.

neighbours is orange. Similarly, node 2 belongs to the orange and green communities. Node 3, however, only belongs to the green community.

Agent and community resource dynamics. We allow agents of the same community to undertake joint ventures by investing in risky projects, thereby forming self-financing groups, similarly to the work by Martines et al (Gonzales Martinez et al., 2021). Each community is assigned a risky project with two possible outcomes and probabilities summing to one, and an agent can only invest in projects assigned to communities that it is a member of. This assumption reflects real-world constraints often faced in developing economies, where investment opportunities may be localized due to: limited information flows, the importance of trust within known social circles, and the geographical specificity of ventures (e.g., community-based agricultural projects) (Dupont and Roy, 2023; Fafchamps, 2003; Uphoff et al., 1998). While a simplification of globalized interconnected settings, it allows us to focus on the impact of local network structures and community-level dynamics on investment opportunities and resultant economic outcomes.

A risky project is generated as follows. First, we randomly draw the probability that the project yields a loss from the uniform distribution $U(\ell, 1 - \ell)$ with parameter $\ell \in [0.30, 0.45]$. Let us call this probability \mathbb{P}_{loss} . The probability that the project yields a gain is therefore $\mathbb{P}_{\text{gain}} = 1 - \mathbb{P}_{\text{loss}}$. Finally, we must determine the actual values of the loss and gain. The loss is drawn once per project from uniform distribution $U(L_{\text{lower}}, L_{\text{upper}})$ and the gain from uniform distribution $U(G_{\text{lower}}, G_{\text{upper}})$, with the following relations: $L_{\text{lower}} < L_{\text{upper}} < 1 < G_{\text{lower}} < G_{\text{upper}}$. Figure 2 illustrates a risky project, including branch probabilities and outcomes.

In addition to community projects, each agent has the option to invest in a safe asset with a guaranteed gain of $G_{\text{safe}} > 1$. However, we still wish for projects to be attractive alternatives to the safe asset, so care is taken in choosing parameter values such that the expected return of each project is at least G_{safe} . Lastly, a project is allowed to take place if and only if the total pooled investment towards that project is at least $\theta \in (0, 1)$ multiplied by the total initial wealth of agents able to contribute towards the project. If this requirement is not met, all invested funds are lost. Otherwise, the project returns the gain with probability \mathbb{P}_{gain} and the loss with probability \mathbb{P}_{loss} .

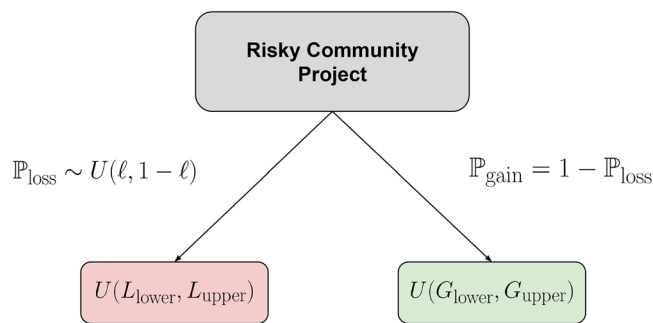


Fig. 2 Diagram of a risky community project. \mathbb{P}_{loss} is independently drawn once for each project, as are the random loss and gain returns in the red and green boxes, respectively.

Portfolio optimisation. Agents must decide how they wish to apportion their capital between the safe asset and whatever risky projects are available to them (De Bruijn and Antonides, 2022; Haushofer and Fehr, 2014). We let the portfolio of agent i be denoted \mathbf{P}_i , which is a vector of weights summing to one. Note that \mathbf{P}_i is at least two-dimensional since every agent can invest in the safe asset and is a member of at least one community with a corresponding project.

In order to select a portfolio, agents must perform portfolio optimisation. This problem is approached using Cumulative Prospect Theory (CPT), which was developed in order to better explain how humans make decisions, particularly under uncertainty. The following theoretical overview of CPT closely parallels that of (Luxenberg et al., 2024)

Agents are assumed to be risk-averse for gains. However, they are also risk-seeking for losses. For a given return x , the prospect theory utility function is given by

$$u^{\text{prosp}}(x) = \begin{cases} u_+(x) & x \geq 0 \\ u_-(x) & \text{otherwise} \end{cases}, \quad (2)$$

where $u_+(x) = 1 - e^{-\gamma_+ x}$ and $u_-(x) = e^{\gamma_- x} - 1$. Here, γ_+ and γ_- are parameters shaping the utility function for gains and losses, respectively, reflecting different risk attitudes. There is the additional requirement that $\gamma_- > \gamma_+ > 0$, which translates to loss aversion. Figure 3 provides a visual representation of Equation

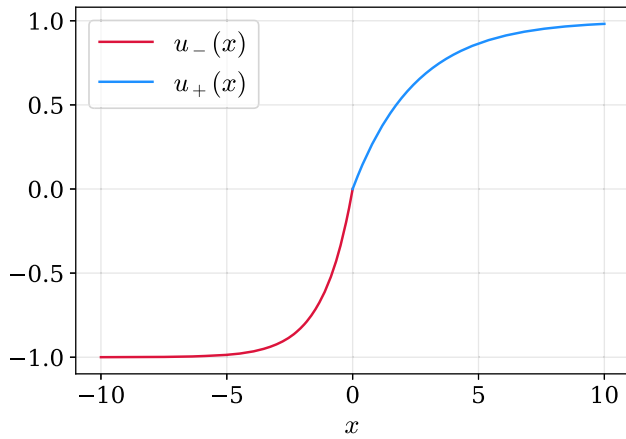


Fig. 3 Prospect theory utility function. Parameters: $\gamma_- = 0.85$ (risk attitude for losses) and $\gamma_+ = 0.4$ (risk attitude for gains).

(2). The steeper exponential curve for losses (in red) can be interpreted as how a small decrease in wealth would decrease the utility more than an equal increase in wealth would increase utility. In other words, losses are felt more heavily than gains.

Next, we define some empirical distribution of returns for n assets. We denote these returns r_1, r_2, \dots, r_N where $r_j \in \mathbb{R}^n$. That is, each r_j is a vector of n realized asset returns. Now consider some portfolio $p \in \mathbb{R}^n$, such that all weights sum to one. The resulting portfolio returns would be $r_1^\top p, r_2^\top p, \dots, r_N^\top p$ where $r_j^\top p \in \mathbb{R}$.

Now let N^- denote the number of negative portfolio returns, and let N^+ denote the number of non-negative portfolio returns, with $N = N^+ + N^-$. After sorting all of the portfolio returns, we can write the following relations:

$$p^\top \rho_1 \leq \dots \leq p^\top \rho_{N^-} < 0 \leq p^\top \rho_{N^-+1} \leq \dots \leq p^\top \rho_N,$$

where ρ_i denotes the i th portfolio return after sorting. In other words, $p^\top \rho_1$ is the biggest loss, and $p^\top \rho_N$ the biggest gain. Consider also the following probability weighing functions, where δ_+ and δ_- are parameters between 0 and 1 that capture how agents subjectively distort objective probabilities:

$$w_+(p) = \frac{p^{\delta_+}}{(p^{\delta_+} + (1-p)^{\delta_+})^{1/\delta_+}}$$

$$w_-(p) = \frac{p^{\delta_-}}{(p^{\delta_-} + (1-p)^{\delta_-})^{1/\delta_-}}$$

Using these, we define the following positive and negative decision weights:

$$\pi'_{+,j} = \begin{cases} w_+\left(\frac{N^+ - j + 1}{N}\right) - w_+\left(\frac{N^+ - j}{N}\right) & j = 1, \dots, N^+ - 1 \\ w_+\left(\frac{1}{N}\right) & j = N^+, \end{cases}$$

$$\pi'_{-,j} = \begin{cases} w_-\left(\frac{N^- - j + 1}{N}\right) - w_-\left(\frac{N^- - j}{N}\right) & j = 1, \dots, N^- - 1 \\ w_-\left(\frac{1}{N}\right) & j = N^- \end{cases}$$

We wish for π'_+ and π'_- to be monotonically increasing, and thereby assign greater weight to extreme outcomes. Although this is the case for most parameter choices, some can lead to the violation of this desirable property. We can ensure that π'_+ and π'_- are non-decreasing by substituting $\pi'_{+,j}$ with $\min(+/\pi)$ for all $j < \text{argmin}(\pi'_+)$, and substituting $\pi'_{-,j}$ with $\min(-/\pi)$ for all $j < \text{argmin}(\pi'_-)$.

In order for π'_+ and π'_- to each have length N instead of N^+ and N^- respectively, we zero-pad both sequences: $\pi_+ =$

$(\mathbf{0}^{N^-}, \pi'_+)$ and $\pi_- = (\mathbf{0}^{N^+}, \pi'_-)$, where $\mathbf{0}^k$ denotes a zero-vector of dimension k .

Let us also define the following two functions operating element-wise on some vector x :

$$\phi_+(x) = \max(0, x)$$

$$\phi_-(x) = -\min(0, x)$$

Using these, we can finally write the CPT utility function as follows:

$$U^{\text{CPT}}(p) = f_{\pi_+}(\phi_+(u_+(\mathbf{R}p))) - f_{\pi_-}(\phi_-(u_-(\mathbf{R}p))), \quad (3)$$

where \mathbf{R} is the matrix of empirical asset returns and f_π is the weighted-ordered-sum function, defined as

$$f_\pi(x) = \sum_{i=1}^N \pi_i x_{(i)},$$

where $x_{(i)}$ denotes the i th smallest element of vector x and π is assumed to be a monotonically increasing probability vector, which we have ensured is the case.

The CPT optimisation problem to be solved by agents is thus to find a portfolio P maximising the CPT utility:

$$\max_P U^{\text{CPT}}(P), \quad (4)$$

which is non-convex as shown by (Luxenberg et al., 2024). Further, Luxenberg et al. (2024) show that CPT utility can be decomposed as a difference of two functions. The first term is a convex function with concave arguments, and the second term is a convex function with convex arguments. This structure allows us to derive maximum CPT utility (Luxenberg et al., 2024).

In practice, we assign each agent its own CPT utility function. The key CPT parameters for each agent are: γ_+ (risk attitude for gains), γ_- (risk attitude for losses), δ_+ (probability weighting for gains), and δ_- (probability weighting for losses). These parameters are drawn from uniform distributions: $\gamma_+ \sim U(\gamma_+^{\text{lower}}, \gamma_+^{\text{upper}})$, $\gamma_- \sim U(\gamma_-^{\text{lower}}, \gamma_-^{\text{upper}})$, $\delta_+ \sim U(\delta_+^{\text{lower}}, \delta_+^{\text{upper}})$, and $\delta_- \sim U(\delta_-^{\text{lower}}, \delta_-^{\text{upper}})$.

Portfolio update and attention mechanism. An agent does not keep the same portfolio throughout the simulation. Agents learn and adjust their strategies over time based on observed payoffs or shocks, leading to dynamic equilibria that reflect real-world fluctuations. Initially, we generate M random returns for each risky project as well as the safe asset (though these are constant), and each agent i bases its initial portfolio $\mathbf{P}_{i,\text{initial}}$ on this data. However, as projects yield returns at each time step, agents can use this information to update their portfolio choices. For example, an agent may decide to reduce the amount invested towards a project that keeps failing and instead redistribute these funds towards other, more promising projects or even the safe asset.

This is done using an attention mechanism as follows. Each agent i is initialised with an independent attention parameter a_i randomly drawn from the uniform distribution $U(0, 1)$. Then, anytime an agent updates its portfolio throughout the simulation, it computes the following weighted sum:

$$\mathbf{P}_i = (1 - a_i) \mathbf{P}_{i,\text{initial}} + a_i \mathbf{P}_{i,\text{observations}}, \quad (5)$$

where $\mathbf{P}_{i,\text{observations}}$ is the optimal portfolio of agent i based on observed project returns since the start of the simulation. Hence, the higher an agent's attention, the more importance it will place on the project returns it observes during the simulation when choosing a portfolio. Portfolio update times are Poisson-distributed with parameter λ , and each agent has

independently sampled update times. Furthermore, we only allow for updates to be carried out after an initial warm-up period of 5 model steps.

Consumption. At each time step t , each agent i must consume a portion $c_{i,t}$ of its current wealth $w_{i,t}$. The amount consumed is determined by the saving propensity parameter $\beta \in (0, 1)$, and is calculated as $c_{i,t} = 1 - \beta \cdot w_{i,t}$. Once capital has been consumed, the remainder is invested according to the agent's current portfolio, and the agent's wealth at the next time step is the return on investments (both safe and risky).

Summary of model parameters. Table 1 summarises key model parameters. Parameters that we wish to vary for the sensitivity analysis have their values listed in square brackets, denoting

uniform intervals. Other parameters have fixed values for all simulations.

Algorithm 1 summarises the key steps of the agent-based model. All parameters are fixed except five, which are listed at the top of the algorithm as input parameter set. After initializing global attributes such as the network, communities, projects, and random initial returns, we initialize individual attributes for $N = 1225$ agents. Then, we perform $K = 100$ model steps during each of which agents consume some capital, invest the rest according to their portfolio, and may also update their portfolio. Before going to the next step, project returns are divided amongst investors to compute each agent's wealth for the next time step. At the end of the simulation, we have an $N \times K$ matrix of agent wealth trajectories and knowledge of agent attributes, project returns, etc. $K = 100$ model steps were found to be sufficient during preliminary experiments for aggregate statistics to stabilize and reach a steady state.

Algorithm 1 **Agent-based model for multi-level poverty trap formation**

Input: parameter set $\{\ell, G_{\text{upper}}, \beta, \theta, \alpha\}$

begin

Generate initial agent wealths, $w_{i,1} \sim \mathcal{N}(10, 1)$, $i = 1, 2, \dots, N$

Construct SDA graph using homophily parameter α

Perform community detection

Determine minimum required investment cost for each community project using θ

Generate a risky project for each community using ℓ, G_{upper}

Generate M initial project returns

for $i = 1, \dots, N$ **do**

 Assign CPT Utility

 Assign attention parameter $a_i \sim U(0, 1)$

 Generate set T_i of portfolio update times

 Compute initial portfolio $\mathbf{P}_i = \mathbf{P}_{i,\text{initial}}$

for $t = 1, \dots, K$ **do**

for $i = 1, \dots, N$ **do**

if $t \geq 5$ and $t \in T_i$ **then**

 Compute $\mathbf{P}_{i,\text{observations}}$

 Perform portfolio update using Equation (5)

 Agent chooses consumption amount $c_{i,t}$ using β

 Agent contributes remainder of wealth towards projects and safe asset according to their portfolio

for each project do

if investment \geq minimum required investment **then**

 Project return = loss with probability ℓ , or gain with probability $1 - \ell$

else

 Project return = 0

 Let \mathbf{R} be the vector of project returns concatenated with G_{safe}

for $i = 1, \dots, N$ **do**

$w_{i,t+1} = (w_{i,t} - c_{i,t})\mathbf{P}_i \cdot \mathbf{R}$

Table 1 Summary of agent-based model parameters.

Parameter	Description	Value(s)
N	Number of agents	1225
K	Number of model steps	100
λ	Poisson parameter for generating update times	10
μ	Average of initial wealth distribution	10
σ	Standard deviation of initial wealth distribution	1
M	Initial number of project returns	2000
L_{lower}	Lower uniform bound for project loss	0.90
L_{upper}	Upper uniform bound for project loss	0.95
G_{lower}	Lower uniform bound for project gain	1.60
G_{upper}	Upper uniform bound for project gain	[1.70, 8.00]
G_{safe}	Safe asset gain	1.10
ℓ	Parameter for generating project loss probabilities	[0.30, 0.45]
θ	Minimum required project investment	[0.05, 0.95]
β	Saving propensity	[0.70, 0.80]
α	Homophily parameter	[2.00, 12.00]
b	Characteristic distance	See "Agent social network"
$\gamma_{+}^{\text{lower}}$	Lower uniform bound for γ_{+} (risk attitude for gains)	5
$\gamma_{+}^{\text{upper}}$	Upper uniform bound for γ_{+} (risk attitude for gains)	30
$\gamma_{-}^{\text{lower}}$	Lower uniform bound for γ_{-} (risk attitude for losses)	31
$\gamma_{-}^{\text{upper}}$	Upper uniform bound for γ_{-} (risk attitude for losses)	70
$\delta_{+}^{\text{lower}}$	Lower uniform bound for δ_{+} (weighting parameter for gains)	0.50
$\delta_{+}^{\text{upper}}$	Upper uniform bound for δ_{+} (weighting parameter for gains)	0.70
$\delta_{-}^{\text{lower}}$	Lower uniform bound for δ_{-} (weighting parameter for losses)	0.71
$\delta_{-}^{\text{upper}}$	Upper uniform bound for δ_{-} (weighting parameter for losses)	0.90

Experiments. Saltelli sampling is used to generate different combinations of parameters (Saltelli, 2002). We choose to have 1024 unique values per parameter, resulting in a total of 7168 parameter combinations. Furthermore, twenty repetitions with different random seeds are carried out for each parameter combination to account for stochastic effects. This approach avoids dynamically altering parameters during simulation runs, instead focusing on pre-generated sample sets for systematic analysis. It is particularly valuable for identifying key drivers of variability in models with high-dimensional parameter spaces.

In addition to studying the resulting wealth trajectories of individuals and communities, we also examine whether a one-time capital injection can help the poorest agents escape poverty—mobility in this context can be understood to mean the capacity of individuals to escape poverty when receiving assets. In order to do this, we randomly select one parameter combination from the All Poor regime and one from the Some Rich regime. Then, for each one, we let our model run for 100 time steps, at which point we give 10 units of capital to the 100 poorest agents. Subsequently, we allow the model to run for an additional 100 steps and examine whether the targeted agents have managed to avoid falling into poverty. 20 repetitions are performed to account for stochastic effects. Lastly, we study the model’s sensitivity to parameters using a global sensitivity analysis method (Bazyleva et al., 2023, 2024) to compute diffusion coordinates and sensitivity indices, shown in Figs. 1 and 2 of Appendix B, respectively. The appendix also contains additional details regarding the method used.

Results

Transition from extreme poverty to extreme inequality. We identify three distinct regimes of model behaviour, representing characteristic steady state outcomes, based on the wealth trajectories of agents across different parameter combinations and repetitions. In the first regime, which we label All Poor, every agent in the population has a lower wealth by the end of the simulation compared to their starting wealth. This regime is

characterised by extreme poverty, and agents are unable to avoid or mitigate financial losses. In the second regime, the Some Rich regime, one or more agents have a final wealth greater than their initial wealth. Most of the time, this wealth accumulation is experienced by only a small minority, thereby giving rise to significant inequality and wealth disparity. Finally, in the third regime, All Rich, every agent has a final wealth greater than their initial wealth. We observe the following proportions of simulations that land in each regime: 12.24% (All Poor), 87.7% (Some Rich), and 0.06% (All Rich). It is worth noting that the All Poor and Some Rich regimes correspond very closely with single and double equilibrium poverty traps, as studied in depth by (Barrett and Carter, 2013) Although formal fitting to specific parametric wealth distribution models was not carried out in this study, we wished to validate the steady-state wealth distributions generated by our model, particularly for the more realistic Some Rich regime. Therefore, we computed the normalized Complementary Cumulative Distribution Function (CCDF) for our steady-state wealth distributions. As illustrated in Fig. 8 (Appendix H), a sample of these simulated CCDFs was compared against an empirical CCDF derived from the UK wealth distribution, with wealth values normalized for appropriate comparison. The figure demonstrates a good alignment between our simulated wealth distributions and the empirical data, suggesting that our model is able to capture key characteristics of real-world wealth inequality, especially in its general shape and tail behaviour.

Our findings show that the evolution of poverty and inequality (state space) is driven by various social, economic, and financial structures (represented by the parameter space of the model) and can explain the observed transition from extreme poverty to pervasive inequality. Figure 4 illustrates the typical parameter ranges giving rise to each of the three regimes. The influential parameters were identified using global sensitivity analysis, as described in Appendix B (see Figs. 1 and 2). The scale of each parameter is adjusted using min-max normalization such that 0 corresponds to the minimum possible value of that parameter,

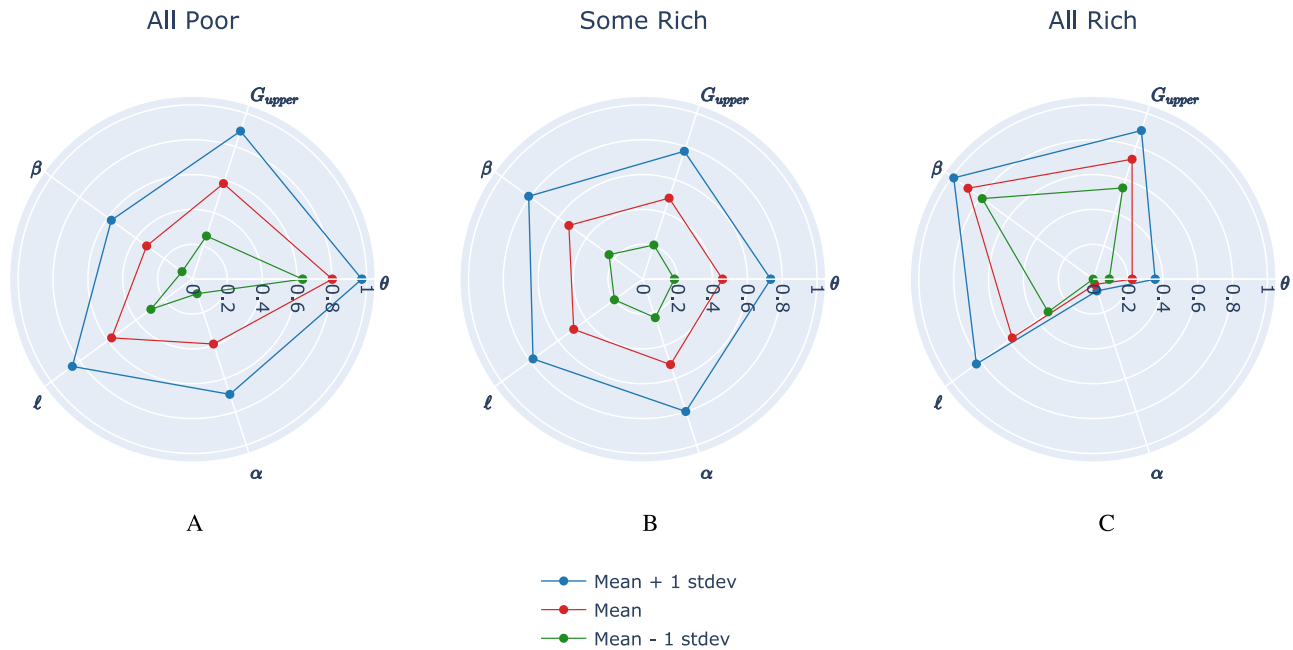


Fig. 4 Average Scaled Parameter Profiles Across Wealth Regimes. Average parameter values (in red) resulting in each of the three distinct regimes: All Poor, Some Rich, All Rich. Key parameters displayed are: θ (minimum project investment threshold), G_{upper} (upper bound for project gain magnitude), β (saving propensity, where a lower value means more consumption), ℓ (probability of project loss), and α (homophily in social network formation, where a higher value means more initial wealth-based segregation). Parameter values have been independently scaled between 0 and 1 based on the ranges defined in Table 1. In blue and green are one standard deviation above and below the average parameter values, respectively. **A** Parameter values giving rise to the All Poor regime. High values of homophily (α) and project cost (θ) are likely to yield this regime. **B** Parameter values giving rise to the Some Rich regime. **C** Parameter values giving rise to the All Rich regime. Homophily and project cost are notably smaller for simulations landing in this regime.

and 1 corresponds to its maximum value (see Table 1 for a summary of model parameters). For each regime, the vertices along the red polygon correspond to average parameter values giving rise to that regime. The analysis of the parameter space reveals three key insights for reducing poverty and inequality.

First, parameter θ , which modulates the minimum project cost required of each community in order for a project to be carried out, plays a significant role in shaping model behaviour. On average, θ is close to 0.8 of its maximum value in the All Poor regime (implying that project costs are high), whereas it is only 0.45 for the Some Rich regime, and 0.2 for the All Rich regime. Clearly, reducing the cost of undertaking community projects is helpful in lowering the barrier to financial well-being and tends to increase the likelihood of agents avoiding poverty. This indicates that reducing frictions related to financial inclusion can increase the level of entrepreneurship, allowing financially constrained individuals to become entrepreneurs and move out of poverty. At the aggregate level, these effects could ultimately boost economic activity, reduce poverty, and potentially increase income equality—as observed in previous empirical studies (Omar and Inaba, 2020; Park and Mercado, 2015).

The parameter ℓ —which determines the probability of projects yielding a loss rather than a gain—does not appear to be different across regimes, while the upper uniform bound for the magnitude of project gains, G_{upper} , is on average significantly higher for the All Rich regime. This implies that increasing project returns is more favourable to agents than increasing the frequency of project success. Second, the saving propensity parameter β also plays an important role in giving rise to the different regimes. In the All Poor regime, β typically takes lower values, indicating that agents consume more and invest less towards risky projects. As pointed out by Banerjee and Duflo, a “main challenge for the poor who try to save is to find safety and a reasonable return”

(Banerjee and Duflo, 2007). Indeed, in the All Poor regime, projects tend to fail more often than not, making it difficult for agents to invest their savings. However, agents in the All Rich regime tend to save more of their wealth in order to invest it towards high-return projects, which tend to be successful. This is in line with traditional perspective on micro and macroeconomic growth, which posits that increased savings, when channelled into investment, would fuel economic growth in households (Steinert et al., 2018). Several empirical studies show that higher savings rates not only directly boost investment but also indirectly increase steady-state output, as rising savings often correlate with rising income, thus stimulating further investment (Do, 2023; Karlan et al., 2014). Lastly, we find that the homophily parameter α must take on particularly low values in order for the All Rich regime to emerge. This translates to lower homophily, implying that there is less initial economic segregation based on wealth, thereby improving the chances of poorer agents to be included in community projects with higher investments and chance of success. By contrast, in the Some Rich regime, homophily tends to be higher, which implies more segregation and leads to greater eventual inequality. The share of high-socioeconomic status (SES) friends among individuals with low SES (low homophily)—known as economic connectedness—is among the strongest predictors of upward income mobility (Chantarat and Barrett, 2012; Chetty et al., 2022; Pena-López et al., 2021; Tóth et al., 2021). A recent study identified that people endowed with high levels of economic and human capital enjoy improved accessibility and networks with a high prevalence of instrumental relations (Pena-López et al., 2021). There is essential inequality in the endowment of social capital, which augments economic inequality. Further, when inequality is socially embedded, traditional re-distributive policies demonstrate limited effectiveness (Pena-López et al., 2021).

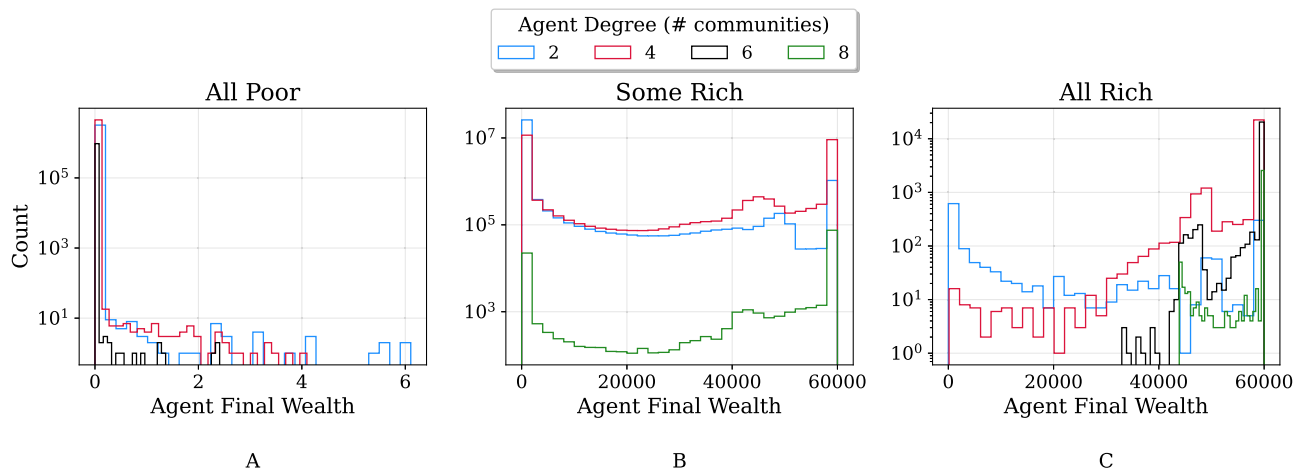


Fig. 5 Effect of access to communities on economic prosperity. **A** Distribution of final agent wealth for All Poor regime simulations. Colours correspond to the number of communities that an agent is part of. In this case, a lower degree seems to be indicative of higher final wealth. **B** Same as sub-figure A but for the Some Rich regime. The dependence of final wealth on agent degree is less evident. **C** Same as sub-figures A and B, but for the All Rich regime. In this case, a higher degree appears to be associated with higher final wealth.

Beyond individual and distributional outcomes, aggregate wealth dynamics differ across the three regimes. The All Poor regime is characterised by a net destruction of aggregate wealth over time since projects consistently fail or yield low returns, and consumption erodes capital. Conversely, the Some Rich regime demonstrates the model’s capacity for aggregate wealth accumulation, even if wealth becomes highly concentrated among a minority of agents. The rare All Rich regime shows robust growth in both individual and aggregate wealth. This capacity for aggregate wealth generation in the Some Rich and All Rich regimes is crucial, as it theoretically provides the resources that could be targeted by redistributive interventions, although the effectiveness of such interventions depends on the underlying systemic factors, as explored in “Asset-based interventions are context dependent”. Nonetheless, Fig. 4 indicates that growth, in terms of total wealth, is most likely when project cost (θ) and homophily (α) are low, and saving propensity (β) is high.

Figure 5 illustrates the relationship between the final wealth of agents in each regime and the number of projects available to the agents (the community degree). Interestingly, although all agents tend towards poverty, it seems favourable in the All Poor regime to be part of fewer communities (and thus investment projects). This can be seen from the tails of the black, red, and blue distributions extending further to the right as community degree decreases, respectively. This story is completely flipped in the All Rich regime, where a higher community degree appears to lead to a higher final wealth. Given that projects, on average, are successful more often in the “All Rich” regime, it indeed makes sense for an agent to diversify their investments. By contrast, because projects fail often in the All Poor regime (for instance, due to higher θ values), agents tend to invest in fewer projects, thereby taking fewer risks and maintaining a larger investment towards the safe asset. This pattern of behaviour is in fact supported by a growing body of literature. For instance, experimental work by Yesuf and Bluffstone with poor households in Ethiopia reveals a high degree of risk aversion and risk-averting behaviour with “perhaps significant implications for long-term poverty” (Yesuf and Bluffstone, 2009). Indeed, farmers are often prone to sub-optimal decision making due to being less willing to undertake projects with high expected returns, which results from near constant exposure to risk factors in their daily work (Rosenzweig and Binswanger, 1993).

Figure 6 displays the aggregated final agent wealth distributions for each regime, as well as corresponding Gini indices (calculated at the end of each simulation at the population level). While agent wealth levels quickly taper off in the All Poor distribution, both of the Some Rich and All Rich distributions have a tail of very rich agents at the far right. Moreover, the All Rich distribution does not have any data points at or near zero since agents have all attained a higher wealth level than their starting point.

In the All Poor and All Rich regimes, Gini coefficients are understandably very low (with a few exceptions in the All Poor regime), since there is little inequality due to all agents being either poor or rich as regime names indicate. However, the intermediate Some Rich regime is characterised by high levels of inequality, with Gini coefficients generally taking on values greater than 0.8. Further details on the relationship between total population wealth and the Gini index for the Some Rich regime may be found in Fig. 5 of Appendix E. The reader may also refer to Fig. 4 of Appendix D to examine the distributions of project returns by regime. While the Gini coefficient provides a useful summary measure of inequality, analysing the full wealth distribution offers further insights. For instance, in representative Some Rich simulations, the top 10% of agents frequently hold well over 80–90% of the total wealth, underscoring the high concentration of wealth suggested by the Gini coefficient. Future work could further explore alternative metrics in order to better capture such nuances.

Overall, the transition from extreme poverty to extreme inequality found here is reminiscent of the global historical trend presented in the Introduction, wherein we argue that many nations have succeeded in reducing poverty over the last decades at the cost of increased inequality.

Emergence of multi-level poverty traps. Similar to our definition of regimes for individual wealth trajectories, we may also define regimes at the community level as follows. In the All Poor regime, each community’s final total wealth is less than its total initial wealth. Meanwhile, in the Some Rich regime, some communities have final total wealth greater than total initial wealth, and all communities have higher final total wealth in the All Rich regime.

After categorizing all model runs into these community-level regimes, we construct another radar plot similar to Fig. 4 in order to examine the parameter ranges giving rise to each regime.

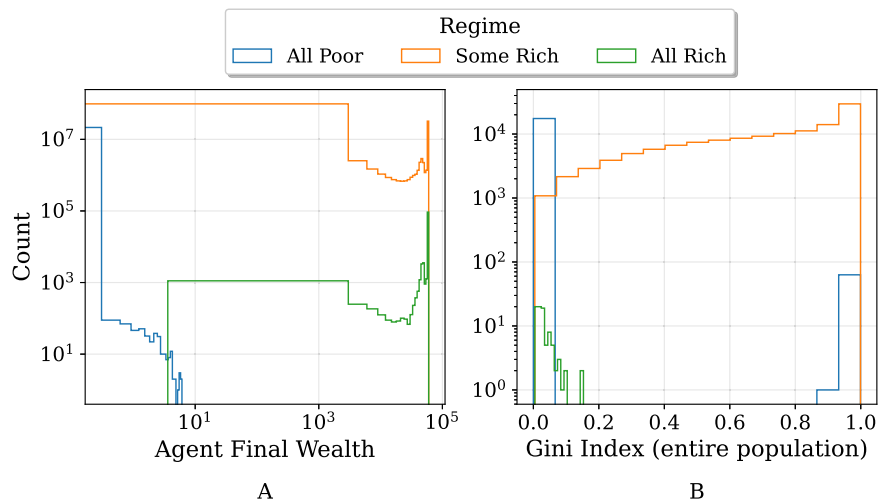


Fig. 6 Comparison of final agent wealth distributions and Gini indices. Data have been aggregated from all simulation runs (7168 parameter combinations \times 20 repetitions each), with each run categorized into the All Poor, Some Rich, or All Rich regime based on final agent wealths. **A** Final wealth of agents in the All Poor regime never surpass small values, whereas agents in the Some Rich regime can attain high levels of wealth. Meanwhile, in the All Rich regime, final wealth levels close to zero are, by definition, impossible. **B** Comparison of Gini indices (at final time step) for populations in different regimes. The All Poor and All Rich regimes tend to exhibit low Gini indices (signifying low inequality), whereas the Some Rich regime exhibits mostly high Gini indices and therefore inequality.

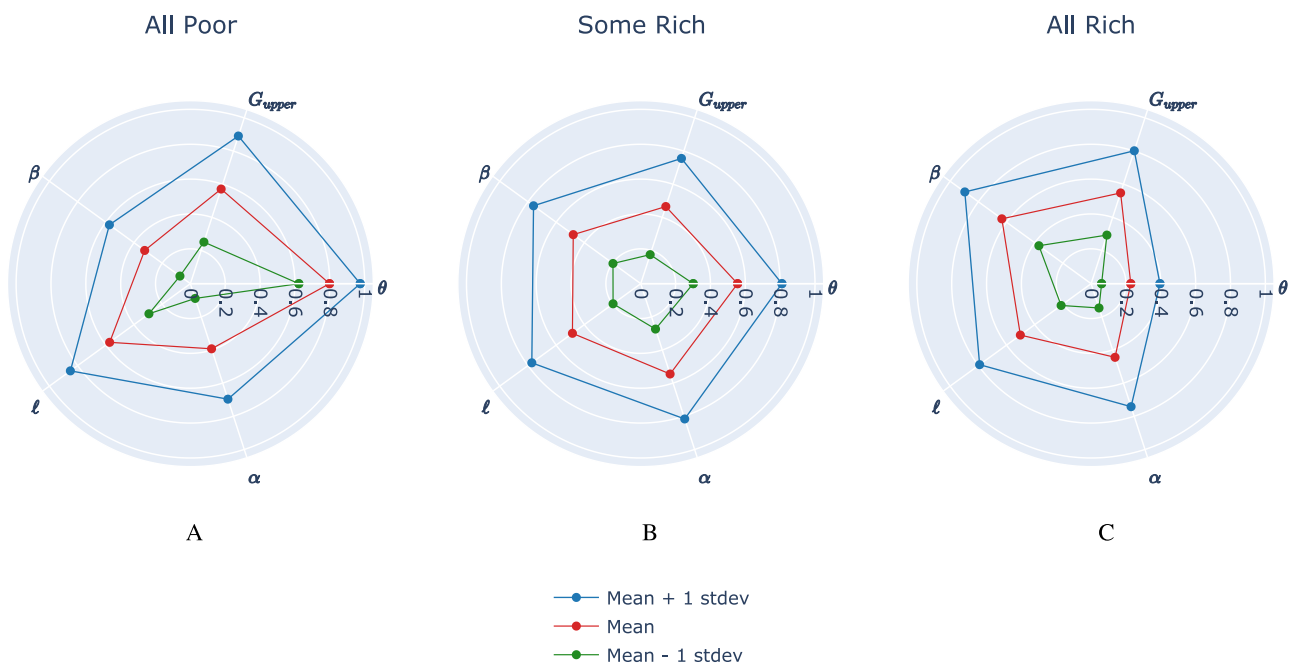


Fig. 7 Average parameter values giving rise to different regimes at the community level. Key parameters displayed are: θ (minimum project investment threshold), G_{upper} (upper bound for project gain magnitude), β (saving propensity, where a lower value means more consumption), ℓ (probability of project loss), and α (homophily in social network formation). Parameter values have been independently scaled between 0 and 1 based on the ranges defined in Table 1. In blue and green are one standard deviation above and below the average parameter values, respectively. **A** High values of project cost (θ) are likely to yield the All Poor regime. **B** Parameter values for the Some Rich regime. **C** Parameter values for the All Rich regime at the community level. The effect of homophily is not as pronounced as it was for the individual-level regimes (see Fig. 4). Higher values of consumption (β) appear more likely for the All Rich regime compared to the Some Rich and All Poor community-level regimes.

Figure 7 displays these results. The proportions of runs for each regime are as follows: 12% (All Poor), 62% (Some Rich), and 26% (All Rich). Interestingly, we observe a much higher proportion of runs landing in the All Rich regime at the community level compared to the individual level regime definition. The All Poor and Some Rich radar plots look quite similar to those of Fig. 4, but the All Rich has striking differences. Although project costs

(θ) still need to take on small values in order for all communities to increase their total wealth over the course of the simulation, we notice that lower homophily (small α) is not as strong of a precondition. Furthermore, simulations giving rise to the All Rich community regime tend to have lower values of both β (saving propensity) and G_{upper} (upper bound for project returns) compared to the All Rich radar plot of Fig. 4.

Table 2 Joint number of simulations landing in individual and community-defined regimes.

Community regime	Individual regime	Count
All poor	All poor	17,537
All poor	Some rich	229
All poor	All rich	0
Some rich	All poor	0
Some rich	Some rich	88,111
Some rich	All rich	0
All rich	All poor	0
All rich	Some rich	37,396
All rich	All rich	87

Interestingly, we identify 229 simulations landing both in the All Poor *community* regime and the Some Rich *individual* regime. This indicates that it is possible for some agents to acquire significant wealth, despite all communities tending towards poverty. Conversely, we also find that as many as 26% of our simulations correspond both to the All Rich community regime and the Some Rich individual regime. This implies that we frequently find societies where all communities have a high total final wealth despite only some agents accumulating wealth. In other words, while the total wealth of communities can significantly increase over the course of a simulation, it can also be the case that many agents remain significantly impoverished. Table 2 summarises the (joint) number of simulations landing in individual and community-defined regimes.

Overall, we conclude that our model gives rise to poverty traps at three distinct levels. One such trap is clearly at the *macroscopic* scale, where the population as a whole tends towards poverty. Interestingly, we also observe a high number of simulations where only some agents become rich and others do not, hinting at high levels of inequality on a global scale. Another set of traps is at the *meso* scale (community). We find simulations where communities tend towards poverty, but have a few rich community members able to avoid poverty. We also have communities tending towards poverty, with all community members also being poor. Finally, there are mechanisms of traps at the *micro* (individual) scale—we find simulations where communities are overall rich, but some community members are stuck in poverty. In fact, as mentioned earlier, as many as 62% of simulations give rise to societies where some communities thrive and others do not. We thus find that inequality and poverty traps can arise at the population level, but also within and between communities. This result is consistent with many empirical findings, in particular work by Barrett et al., where the authors argue that poverty traps can appear not only at the individual or household level, but also at community, regional, and even national scales (Barrett and Carter, 2013). This multilevel poverty trap approach offers a distinct advantage by enabling an effective allocation of resources for interventions at the most appropriate levels. By understanding how interventions at different levels can have cascading effects, we can ensure targeted and coordinated efforts. Unlike relying on ad hoc decisions or anecdotal experiences, stylized modelling facilitates reasoning through various poverty alleviation strategies and outcomes, aiding practitioners in making informed decisions.

Agent heterogeneity, information poverty, and sub-optimal diversification. In order to isolate and examine the effect of individual differences (attention and CPT utility parameters) between agents, we group agents according to the set of projects that they have access to. For instance, one group of agents could consist of those who have access to community projects 2 and 3,

and another group of agents could contain those who have access to community projects 2, 3, and 5. These would be distinct groups, even though all agents in the second group have access to all projects accessible to agents in the first group.

Let us first pay attention to the top row of Fig. 8, where we highlight the impact of agent attention on the magnitude of economic shocks (defined as decreases in wealth between subsequent time steps) for agents with access to community projects 9 and 12 in a given simulation. The leftmost figure reveals a positive correlation between attention and the magnitude of economic shocks that agents experience—higher attention translates to experiencing more significant shocks. In order to explain this outcome, we want to better understand how projects 9 and 12 are performing over the course of the simulation. The middle figure reveals that project 9 fails consistently for the first 25 steps or so. Agents with high attention will notice this and invest more capital towards project 12 when updating their portfolio. On the other hand, agents with lower attention will ignore these empirical returns and favor their initial portfolio. The rightmost figure demonstrates this clearly, where we notice that lower attention results in investing a larger proportion of capital towards project 9 by the end of the simulation. This investment, although initially dubious, ends up benefiting agents with low attention since they tend to have more diversified portfolio and are thus able to better absorb economic shocks compared to agents with high attention.

Now let us take a look at the second row of Fig. 8 for an example of how CPT utility parameter γ_+ (which models an agent’s risk attitude towards potential gains) can impact the final wealth of agents with access to projects 0 and 6. The leftmost figure shows a steep transition where agents with $\gamma_+ < 15$ are extremely rich by the end of the simulation whereas agents with a value for this parameter greater than 15 converge towards a state of poverty. The central figure shows that both projects 0 and 6 yield nonzero returns throughout the entire simulation. In fact, these empirical returns agree very closely with what we would expect, unlike the previous example, where project 9 failed for numerous subsequent steps. Therefore, in this scenario, attention plays a rather insignificant role when agents update their portfolios. The difference in final wealth is instead attributable to γ_+ , and we can observe from the rightmost figure that a lower γ_+ value translates to a higher proportion of wealth being invested in the safe asset. In other words, for this particular group of agents, it is more beneficial to take less risk by investing in the safe asset, presumably due to the risky and volatile nature of the project returns. Although not represented in this rightmost figure, we may even expect that too little risk-taking could end up being harmful (for instance, if the safe investment proportion was as high as >0.70), since agents would not end up benefiting from the high returns of the projects as much.

To conclude this section, we remark that examples such as the ones of Fig. 8 abound throughout the many simulations that were carried out. We have restricted the present analysis in order to showcase the complex and rich behaviour that can be observed using our model as a result of heterogeneity in the agent population. Importantly, whether attention and γ_+ (risk attitude towards gains) are related to economic shocks or final wealth depends on numerous factors, such as which group of agents is considered, but also on which model parameters and regimes are being examined. This outcome highlights the importance of considering traps at multiple levels, since behaviour can vary drastically between different groups or communities of agents even when they are part of the same larger population. Two key insights emerge from our analysis. First, a crucial link may be established between the attention mechanism in our model and the notion of “information poverty” in the literature, which is a

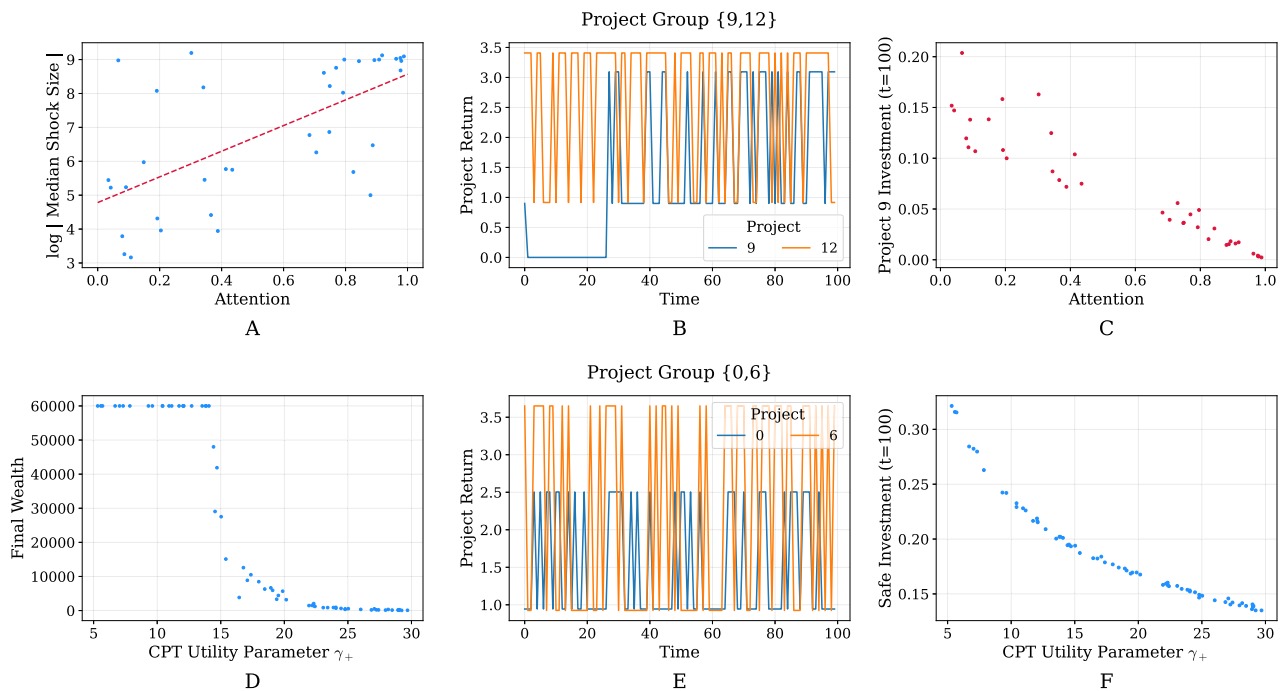


Fig. 8 Impact of agent attention and risk attitude towards gains on shocks, final wealth, and investment strategy. Sub-figures A through C show results from a single, hand-picked simulation run that resulted in the Some Rich regime. Sub-figures D through E similarly show results from a different single, hand-picked simulation run resulting in the Some Rich regime. **A** Impact of agent attention on the median shock size experienced for agents with access to the same two community projects (projects 9 and 12). Agents with higher attention experienced economic shocks of greater magnitude. **B** Returns of community projects 9 and 12. **C** Impact of agent attention on the proportion invested in community project 9 by the end of the simulation. **D** Impact of the γ_+ CPT utility parameter (risk attitude for gains, influencing investment choices) on agent final wealth for agents with access to the same two community projects (projects 0 and 6). **E** Returns of community projects 0 and 6. **F** Impact of γ_+ CPT utility parameter (risk attitude for gains) on the proportion invested into the safe asset by the end of the simulation.

situation in which individuals or communities do not have the necessary skills or knowledge to obtain information, as well as to interpret and apply it correctly to a particular problem (Marcella and Chowdhury, 2020). Agents with higher attention pay greater attention to empirical observations, implying that they have better access to information. Similarly, in real-world settings, individuals or communities with greater access to sources of information, such as education or the Internet, are less likely to experience information scarcity, thereby empowering them to make more informed decisions and respond appropriately to different scenarios that may arise. While this analogy is not perfect, it does underscore the importance of what information agents have access to, and crucially, what they decide to do with it: even minor differences in information access and information use can lead to drastically different outcomes for different agents. Second, we show that suboptimal diversification due to information acquisition, particularly in the context of investments or economic activities, can indeed have an impact on poverty. Suboptimal diversification can lead to increased volatility in income (Capuano and Ramsay, 2011; Van Nieuwerburgh and Veldkamp, 2010). Previous research shows that risk-averse individuals tend to focus more on past outcomes (high attention), especially negative ones, when making decisions under uncertainty (Dillenberger and Rozen, 2015; Li et al., 2023). This heightened attention to prior losses or adverse experiences can increase their sensitivity to risk and drive more cautious choices, leading to suboptimal diversification. For example, if a household relies solely on one unstable source of income, such as seasonal agricultural work, it may face periods of scarcity when that income dries (Michler and Josephson, 2017). This can push them into poverty or exacerbate their existing poverty. In times of

economic downturns, natural disasters, or other crises, families with sub-optimal diversification may lack resilience to bounce back, potentially slipping into poverty or facing prolonged periods of financial hardship (Antonelli et al., 2022). At the community level, a lack of diversification can lead to widespread unemployment and poverty if, for example, a community heavily relies on a single industry that faces a downturn or technological disruption.

Asset-based interventions are context dependent. Lastly, we turn our attention to the results of intervention experiments, as described in “Experiments”. Figure 9 displays a sample of three out of 20 interventions that were performed for a choice of parameters corresponding to the All Poor (top row) and Some Rich (bottom row) regimes. At $t = 100$, the poorest 100 agents were identified and received a wealth injection of 10 units, after which the model was allowed to run for another 100 time steps. The trajectories displayed below are those corresponding to this subset of agents only.

For the All Poor regime, we find that in all 20 repetitions of our experiments, the targeted agents all fall back into poverty with final wealth approaching zero ($<10^{-8}$) by $t = 200$. This indicates that interventions that do not change the underlying landscape of factors (in this case the parameter space) contributing to poverty will not favour impoverished agents in the long term. By contrast, for the Some Rich regime, approximately 32% of targeted agents manage to escape poverty by $t = 200$, with a standard deviation of 13%. Note that here we define the poverty line as being set by the wealth of the wealthiest agent among the 100 poorest agents at the end of the experiment. This outcome implies that when the underlying parameter space is conducive to upward mobility, a

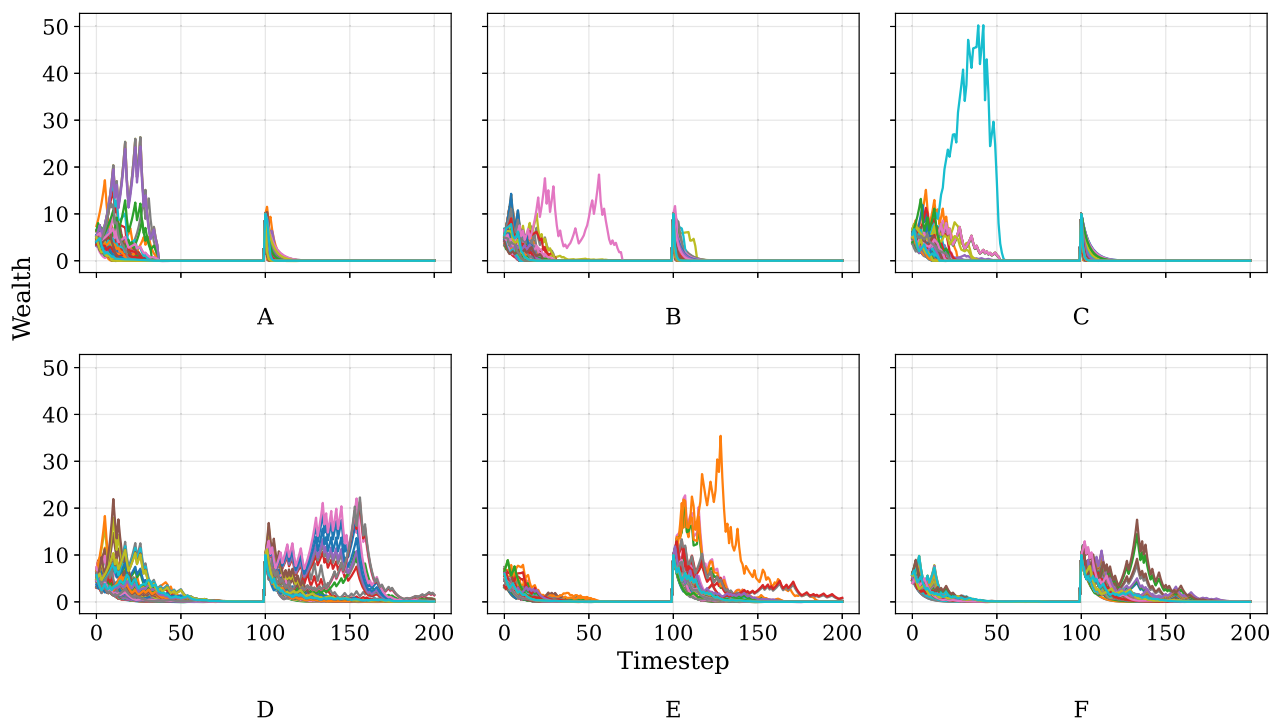


Fig. 9 Examples of wealth trajectories for the 100 poorest agents following a one-time capital injection after 100 time steps. **A-F** display results from a distinct, single simulation run chosen to illustrate typical outcomes. Three example runs are shown for interventions in the All Poor regime (top row, **A-C**) and three for the Some Rich regime (bottom row, **D-F**). Interventions were performed for a specific parameter combination representative of each regime, with 20 Monte Carlo repetitions conducted for each; these plots only show illustrative individual repetitions.

capital injection can prove to be successful at helping agents to escape from poverty.

This contrast in effectiveness between regimes suggests that interventions targeting systemic parameters (like θ , the minimum project investment threshold, which represents a barrier to financial inclusion/entrepreneurship) can have more profound and lasting effects than one-time capital injections when the underlying system dynamics are unfavourable. Lowering θ fundamentally alters the “rules of the game”, making projects systemically more accessible and viable for all agents within that community over the long term. A cash injection, while potentially enabling an agent to meet the θ threshold for one or a few investment cycles, does not change the high cost or low viability of projects if θ remains high and other adverse conditions persist. Thus, without systemic changes to opportunity structures, individual boosts may offer only temporary relief in the All Poor regime. It is important to note that our cash transfer intervention was modelled with a fixed magnitude (10 units) and a fixed number of beneficiaries (100 poorest agents); exploring variations in these parameters, such as larger transfers or different targeting strategies, could yield different insights and is a subject for future investigation.

The efficacy of capital injection, or cash transfers, as a means of combating poverty remains a controversial topic even today. A recent article by Miao and Li examining the impact of cash transfers towards rural households in China casts doubt on whether this form of intervention necessarily guarantees positive long-term effects (Miao and Li, 2023). Although the cash transfers enabled poor rural residents to spend more on healthcare, they also led to a decrease in labour supply and no change in education spending, despite education being a significant factor in long-term poverty mitigation. Our results offer a similar point of view, especially for the All Poor regime, in the sense that cash transfers run the risk of only providing short-

term relief with no concrete long-term outcomes as a result of not addressing underlying structural and systemic issues (such as the θ parameter, which prohibits project execution when too high).

Discussion

The transition from extreme poverty to extreme inequality highlights the importance of addressing various interconnected challenges to achieve Sustainable Development Goals effectively. Here, we present a pathway that outlines steps toward achieving sustainable development while addressing both poverty and inequality.

Inequality drives persistent poverty. In the Introduction, we made two observations. First, although extreme poverty has decreased in the last three decades, recently about 97 million more people are living on less than \$1.90 a day due to the COVID-19 pandemic, increasing the global poverty rate from 7.8 to 9.1 percent. Globally, three to four years of progress toward ending extreme poverty are estimated to have been lost. Second, between-country inequality is estimated to be increasing for the first time in a generation. Emerging evidence even shows that within countries, inequality may also have worsened. The key question is: are we stuck between extreme poverty and extreme inequality? Why is progress so vulnerable to shocks?

An examination of the parameter space reveals that our model captures two extreme regimes: one characterized by poverty and the other prosperity. Both regimes have low inequality. An interesting third regime emerges where some agents are trapped in poverty while others can escape poverty and accumulate wealth. An analysis of the Gini coefficient of simulations in this regime shows that inequality is negatively correlated with economic development (see Fig. 5 in Appendix F)—a fact well established in empirical literature (Alfani and Schifano, 2021).

High inequality inhibits social mobility perpetuated through increased segregation in social connections, high cost to access financial services, sub-optimal diversification and low returns to investment. Although our analysis primarily focuses on agents escaping poverty (a form of upward absolute mobility), a different approach to measuring social mobility could involve measuring rank dynamics, such as the “stickiness” of an agent’s or community’s position within the wealth distribution over time, as suggested by recent literature (Stojkoski, 2024). Our findings in “Asset-based interventions are context dependent,” where some interventions allow agents to escape poverty in the Some Rich regime, provide initial insights into factors promoting such upward movements. Future work could explicitly quantify these transition dynamics.

While there are clear pathways in the parameter space to move from the All Poor regime to All Rich via the Some Rich regime, our model indicates two key challenges. First, progress in poverty reduction is vulnerable to shocks due to the presence of critical transitions, wherein even a small change in the parameter space can lead to regime shifts. Many poverty reduction programs rely on sustained economic growth and stability of the parameter space. Economic shocks, such as recessions, currency devaluations, or commodity price fluctuations, can disrupt growth trajectories and undermine the gains made in poverty alleviation. Lack of savings, access to credit, or insurance can magnify the impact of such shocks, leading to deeper and more prolonged periods of poverty. In an interconnected world, shocks in one part of the world can quickly ripple through to other regions. We show that in a multi-dimensional parameter space, there are multiple pathways through which shocks can cause regime shifts in the economic landscape. Second, in unequal societies, powerful individuals or groups may engage in rent-seeking behaviour, using their influence to capture economic rents or extract resources from the rest of society without creating corresponding value. This can exacerbate inequality by diverting resources away from productive uses and concentrating wealth in the hands of a few. The intersection of inequality and political economy is a rich area of study that examines how economic policies, institutions, and power dynamics shape and perpetuate unequal distributions of resources, opportunities, and outcomes within societies. While we do not model such a feedback explicitly, the analysis of the homophily parameter in our model is an excellent proxy for the political economy and validates this finding. Finally, horizontal inequalities (see Fig. 7 in Appendix G) refer to disparities among distinct identity groups and communities, such as racial divides between blacks and whites, gender disparities between women and men, religious differences like Muslims and Hindus, or ethnic tensions like Hutus and Tutsis, among various other examples. These inequalities are inherently unjust and often endure over time. Our results indicate that horizontal inequality is significantly higher than vertical inequality. This demonstrates that an individual’s position in the distribution is mainly determined by the community (group) they belong. Deprived groups face formidable obstacles, including limited financial resources for investing in assets and education for their children, as well as restricted access to loans. Additionally, social networks typically remain segregated within these groups, resulting in fewer beneficial connections for individuals from disadvantaged backgrounds seeking access to quality education or employment opportunities (Roy et al., 2018a). The Sustainable Development Goals acknowledge the imperative of addressing inequalities not only at the individual level but also among different identity groups.

As discussed next, addressing these vulnerabilities requires a multi-dimensional approach that includes building resilience at individual, community, and institutional levels, implementing

adaptive policies, and fostering sustainable development strategies that are less susceptible to external shocks.

Pathways to poverty alleviation and social mobility. What types of poverty alleviation interventions are possible, and how do they depend on the different pathways of inequality and poverty found at different societal scales? In recent work (Lade et al., 2017), using a resilience-thinking lens, describe three different types of interventions that assume the existence of a barrier to escaping poverty.

The first type of interventions involves capital inputs to the poor state in order to “push it” over existing barriers, for example, cash transfers or increased agricultural inputs. Although effective in some cases, it is important to note that such interventions do not lead to a corresponding systemic change and therefore may not yield substantial long-term benefits. Our intervention experiments on artificial societies in the All Poor regime appear to replicate this particular phenomenon: mere capital injection did not alter the fundamental unfavourability of the investment landscape (e.g., high project cost or poor project returns) or harmful behavioural patterns (e.g., high consumption, low saving propensity) that defined the All Poor trap. While a cash transfer might enable an agent to overcome project costs temporarily, if this cost itself isn’t reduced or if project viability does not improve, the agent is likely to fall back into poverty after some time. Furthermore, as shown in our results on the impact of heterogeneity in the agent population (in the form of attention and utility parameters), additional factors such as access to information and education can be important in improving an individual’s chances of escaping poverty. Furthermore, money transfer programs can sometimes encounter moral hazard issues. Moral hazard refers to the risk that individuals or organisations might change their behaviour in response to the assurance of financial assistance, leading to unintended consequences.

The second type of intervention consists of lowering the barrier by changing the parameters that otherwise serve as key mechanisms in inducing and reinforcing poverty. Although reducing the risk aversion of individuals and communities such that they undertake riskier ventures is technically possible in our modelling framework, in practice, this may not be the case. Individuals may not be willing to accept higher levels of risk when already facing the obstacles of extreme poverty or the near prospect of them. By contrast, we find that improving financial inclusion and access to markets or facilitating the adoption of climate-smart farming practices could realistically help lower difficulties associated with escaping the poverty cycle at the individual and community levels. Further, we show that promoting diversity can also lead to lowering the barrier. Homophily can reinforce existing inequalities by perpetuating patterns of advantage and disadvantage within social networks. For example, individuals from privileged backgrounds may have greater access to social networks with valuable connections and opportunities, while those from marginalized or disadvantaged backgrounds may face barriers to entry or exclusion from such networks. As a result, inequalities in access to resources, opportunities, and social capital can become entrenched and perpetuated across generations. Overall, diversity and access to technology empower individuals, strengthen communities, and foster inclusive and sustainable development. By leveraging technology effectively, policymakers, businesses, and civil society organisations can advance efforts to reduce poverty and inequality and build a more equitable and prosperous future for all.

Lastly, the third type of interventions involves major systemic transformations that bring about fundamentally different poverty and inequality landscapes. An example of such an intervention

would be to provide novel developmental pathways to communities while still adhering to key tenets of social and environmental justice (Lade et al., 2017). One of the key findings of our model is the emergence of sub-optimal diversification. Addressing sub-optimal diversification often involves systemic transformation and broader strategies aimed at innovations (e.g., innovative agricultural practices), improving access to education and skills training (e.g., new challenge-based learning), fostering entrepreneurship (e.g., microloans, savings groups), and building resilience to external shocks. These measures can help break the cycle of poverty and create more sustainable livelihoods. Although this type of intervention is arguably the most challenging to implement, it has the potential to yield the greatest benefit since it encourages a reconfiguration of the relationships between agents, their environment, and economic practices, which can all contribute to persistent poverty if not appropriately addressed. For example, negative consequences could quickly arise from asset input into a community with prevalent gender inequality: while such an intervention could appear successful in terms of aggregate asset ownership, gender inequalities may intensify if women's corresponding access is not ensured in a culturally sensitive manner (Simon, 2019).

Overall, we find that the complex and multi-level nature of poverty traps makes it especially difficult to generalise findings across different contexts: a one-size-fits-all policy or intervention simply does not exist. In a recent paper, Radosavljevic et al. discuss the existence of fractal poverty traps wherein “multiple low-level equilibria exist on different levels at the same time and self-reinforce through cross-level feedbacks” (Radosavljevic et al., 2021). These cross-level dynamics are particularly crucial to understand because interventions at the household level, for instance, could have unforeseen and damaging consequences at the community level. This concern points to the importance of integrating information from multiple levels to better understand and quantify how multi-level interactions contribute to poverty trap mechanisms. The metric we employ to define and track progress across levels is also of critical importance. For instance, although shown to be a poor indicator of quality of life, the gross domestic product remains one of the most commonly used metrics to compare national economies. Integrating alternative indicators such as the Multidimensional Poverty Index (UNDP, 2023) or the Green Growth Index, for instance, could provide a more comprehensive overview of how countries, cities, and communities are performing in terms of poverty reduction and sustainable development targets. Moreover, only looking at the global Gini coefficient for income per capita disguises the glaring inequalities that we have argued still exist at regional and country levels. In fact, it is easy to show that a unimodal income distribution, such as the one we see today at the global level, may be constructed by summing many bimodal distributions exhibiting inequality (see Appendix F and Fig. 6).

Data availability

This study does not use any real-world data; all results are generated from simulation experiments. The complete set of simulation outputs generated and analysed during the current study are available in the Figshare repository: <https://figshare.com/s/23efbca44d7b28ac0340>. The simulation model, along with all scripts necessary to reproduce the results and generate all figures presented in the paper, is available as open-source code at <https://github.com/charlesaugdupont/poverty-trap>.

Code availability

This study does not use any real-world data; all results are generated from simulation experiments. The complete set of

simulation outputs generated and analysed during the current study are available in the Figshare repository: <https://figshare.com/s/23efbca44d7b28ac0340>. The simulation model, along with all scripts necessary to reproduce the results and generate all figures presented in the paper, is available as open-source code at <https://github.com/charlesaugdupont/poverty-trap>.

Received: 21 February 2025; Accepted: 9 October 2025;

Published online: 19 November 2025

References

- Adato M, Carter MR, May, J (2013) Exploring poverty traps and social exclusion in South Africa using qualitative and quantitative data. In: Understanding and reducing persistent poverty in Africa. Routledge, p 60–81
- Aghion P, Bolton P (1992) Distribution and growth in models of imperfect capital markets. *Eur Econ Rev* 36(2-3):603–611
- Aghion P, Caroli E, Garcia-Penalosa C (1999) Inequality and economic growth: the perspective of the new growth theories. *J Econ Lit* 37(4):1615–1660
- Alfani G, Schifano S (2021) Wealth inequality in the long run. In: How was life? Volume II: new perspectives on well-being and global inequality since 1820. OECD Publishing, Paris, France, p 103–123
- Antonelli C, Coromaldi M, Pallante G (2022) Crop and income diversification for rural adaptation: Insights from Ugandan panel data. *Ecol Econ* 195:107390
- Azariadis C, Stachurski J (2005) Poverty traps. *Handb Econ growth* 1:295–384
- Balboni C, Bandiera O, Burgess R, Ghatak M, Heil A (2022) Why do people stay poor? *Q J Econ* 137(2):785–844
- Banerjee A, Yakovenko VM (2010) Universal patterns of inequality. *N. J Phys* 12(7):075032
- Banerjee AV, Duflo E (2007) The economic lives of the poor. *J Econ Perspect* 21(1):141–168
- Banerjee AV, Newman AF (1991) Risk-bearing and the theory of income distribution. *Rev Econ Stud* 58(2):211–235
- Barrett CB, Carter MR (2013) The economics of poverty traps and persistent poverty: empirical and policy implications. *J Dev Stud* 49(7):976–990
- Barrett CB, Swallow BM (2006) Fractal poverty traps. *World Dev* 34(1):1–15
- Barrett CB, Garg T, McBride L (2016) Well-being dynamics and poverty traps. *Annu Rev Resour Econ* 8(1):303–327
- Barrett CB, Carter MR, Chavas JP (2018) The Economics of Poverty Traps. Number barr-3 in NBER Books. National Bureau of Economic Research, Inc
- Bazyleva V, Garibay VM, Roy D (2023) Global sensitivity analysis using polynomial chaos expansion on the grassmann manifold. In: J. Mikyška J, de Mulatier C, Paszynski M, Krzhizhanovskaya VV, Dongarra JJ, Sloat PM (eds) Computational science – ICCS 2023. Cham. Springer Nature Switzerland, p 583–597
- Bazyleva V, Garibay VM, Roy D (2024) Trajectory-based global sensitivity analysis in multiscale models. *Sci Rep* 14(1):13902
- Botta A, Caverzasi E, Russo A, Gallegati M, Stiglitz JE (2021) Inequality and finance in a rent economy. *J Econ Behav Organ* 183:998–1029
- Caiani A, Russo A, Gallegati M (2019) Does inequality hamper innovation and growth? an ab-sfc analysis. *J Evol Econ* 29(1):177–228
- Capuano A, Ramsay I (2011) What causes suboptimal financial behaviour? an exploration of financial literacy, social influences and behavioural economics. In: An exploration of financial literacy, social influences and behavioural economics (March 23, 2011). U of Melbourne Legal Studies Research Paper, p 540
- Chakraborti A, Chakraborti BK (2000) Statistical mechanics of money: how saving propensity affects its distribution. *Eur Phys J B Condens Matter Complex Syst* 17:167–170
- Chancel L, Piketty T, Saez E, Zucman G et al. (2021) World inequality report 2022. Research report, World Inequality Lab, Paris, France. Foreword by Abhijit Banerjee and Esther Duflo
- Chantarat S, Barrett CB (2012) Social network capital, economic mobility and poverty traps. *J Econ Inequal* 10:299–342
- Chetty R, Jackson MO, Kuchler T, Stroebel J, Hendren N, Fluegge RB (2022) Social capital I: measurement and associations with economic mobility. *Nature* 608(7921):108–121
- Chong A, Gradstein M (2007) Inequality and institutions. *Rev Econ Stat* 89(3):454–465
- Cordasco G, Gargano L (2011) Community detection via semi-synchronous label propagation algorithms. p 1–8
- De Bruijn E-J, Antonides G (2022) Poverty and economic decision making: a review of scarcity theory. *Theory Decis* 92(1):5–37
- Dillenberger D, Rozen K (2015) History-dependent risk attitude. *J Econ Theory* 157:445–477

- Do MH (2023) The role of savings and income diversification in households' resilience strategies: evidence from rural Vietnam. *Soc Indic Res* 168(1):353–388
- Dupont C, Roy D (2023) Manifold analysis for high-dimensional socio-environmental surveys. In: International conference on computational science. Springer, p 25–39
- Fafchamps M (2003) Market institutions in Sub-Saharan Africa: theory and evidence
- Fierro LE, Giri F, Russo A (2023) Inequality-constrained monetary policy in a financialized economy. *J Econ Behav Organ* 216:366–385
- Galor O, Zeira J (1993) Income distribution and macroeconomics. *Rev Econ Stud* 60(1):35–52
- Goderis B, Malone SW (2011) Natural resource booms and inequality: theory and evidence. *Scand J Econ* 113(2):388–417
- Gonzales Martinez R, D'Espallier B, Mersland R (2021) Bifurcations in business profitability: an agent-based simulation of homophily in self-financing groups. *J Bus Res* 129:495–514
- World Bank and the Development Research Center of the State Council (2022) The People's Republic of China. Four decades of poverty reduction in china: Drivers, insights for the world, and the way ahead. Technical report, World Bank, Washington, DC
- Hasell J, Rohenkohl B, Arriagada P, Ortiz-Ospina E, Roser M (2022) Poverty. *Our World in Data*. <https://ourworldindata.org/poverty>
- Haushofer J, Fehr E (2014) On the psychology of poverty. *Science* 344(6186):862–867
- Jalan J, Ravallion M (2002) Geographic poverty traps? a micro model of consumption growth in rural china. *J Appl Econ* 17(4):329–346
- Karlan D, Ratan AL, Zinman J (2014) Savings by and for the poor: a research review and agenda. *Rev Income Wealth* 60(1):36–78
- Kraay A, McKenzie D (2014) Do poverty traps exist? assessing the evidence. *J Econ Perspect* 28(3):127–148
- Lade SJ, Haider LJ, Engström G, Schlüter M (2017) Resilience offers escape from trapped thinking on poverty alleviation. *Sci Adv* 3(5):e1603043
- Li W, Chen S, Xiao Z, Li D, Lv C, Zhang S (2023) Risk aversion in risk-taking tasks: combined effects of feedback attributes and cognitive reflection ability. *Brain Behav* 13(5):e2957
- Luxenberg E, Schiele P, Boyd S (2024) Portfolio optimization with cumulative prospect theory utility via convex optimization. *Comput Econ* 25(1):555–574
- Marcella R, Chowdhury G (2020) Eradicating information poverty: an agenda for research. *J Librariansh Inf Sci* 52(2):366–381
- McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: homophily in social networks. *Annu Rev Socio* 27(1):415–444
- Mengesha I, Roy D (2025) Evolutionary game selection leads to emergent inequality. In: International conference on computational science. Springer, p 284–297
- Miao Y, Li Z (2023) The poverty alleviation effect of transfer payments: evidence from China. *Human Soc Sci Commun* 10(1):910
- Michler JD, Josephson AL (2017) To specialize or diversify: agricultural diversity and poverty dynamics in Ethiopia. *World Dev* 89:214–226
- Méjean A, Collins-Sowah P, Guivarch C, Piontek F, Soergel B, Taconet N (2024) Climate change impacts increase economic inequality: evidence from a systematic literature review. *Environ Res Lett* 19:043003
- Mutlu A, Roy D, Filatova T (2023) Capitalized value of evolving flood risks discount and nature-based solution premiums on property prices. *Ecol Econ* 205:107682
- Omar MA, Inaba K (2020) Does financial inclusion reduce poverty and income inequality in developing countries? A panel data analysis. *J Econ Struct* 9(1):37
- Palagi E, Napoletano M, Roventini A, Gaffard J-L (2023) An agent-based model of trickle-up growth and income inequality. *Econ Model* 129:106535
- Park CY, Mercado R (2015) Financial inclusion, poverty, and income inequality in developing Asia. Asian Development Bank Economics Working Paper Series, 63(1)
- Pena-López A, Rungo P, Sánchez-Santos JM (2021) Inequality and individuals' social networks: the other face of social capital. *Camb J Econ* 45(4):675–694
- UNDP (2023) 2023 global multidimensional poverty index (MPI). UNDP (United Nations Development Programme)
- Qureshi Z (2023) Rising inequality: a major issue of our time | Brookings—brookings.edu. <https://www.brookings.edu/articles/rising-inequality-a-major-issue-of-our-time/>. [Accessed 16-08-2023]
- Radicchi F, Castellano C, Cecconi F, Loreto V, Parisi D (2004) Defining and identifying communities in networks. *Proc Natl Acad Sci USA* 101(9):2658–2663
- Radosavljevic S, Haider LJ, Lade SJ, Schlüter M (2021) Implications of poverty traps across levels. *World Dev* 144:105437
- Rosenzweig MR, Binswanger HP (1993) Wealth, weather risk and the composition and profitability of agricultural investments. *Econ J* 103(416):56–78
- Roy D, Lees MH, Pfeffer K, Sloot PM (2018) Spatial segregation, inequality, and opportunity bias in the slums of Bengaluru. *Cities* 74:269–276
- Roy D, Palavalli B, Menon N, King R, Pfeffer K, Lees M (2018) Survey-based socio-economic data from slums in Bangalore, India. *Sci Data* 5(1):1–9
- Saltelli A (2002) Making best use of model evaluations to compute sensitivity indices. *Comput Phys Commun* 145:280–297
- Santos P, Barrett CB (2018) Heterogeneous wealth dynamics on the roles of risk and ability. In: The economics of poverty traps. University of Chicago Press
- Simon CA (2019) The effect of cash-based interventions on gender outcomes in development and humanitarian settings. UN Women Discussion Papers, UN
- Steinert JL, Zenker J, Filipiak U, Movsisyan A, Cluver LD, Shenderovich Y (2018) Do saving promotion interventions increase household savings, consumption, and investments in sub-Saharan Africa? a systematic review and meta-analysis. *World Dev* 104:238–256
- Stiglitz JE (1973) Education and inequality. *Ann Am Acad Political Soc Sci* 409(1):135–145
- Stojkoski V (2024) Measures of physical mixing evaluate the economic mobility of the typical individual. *Chaos Solitons Fractals* 180:114567
- Talaga S, Nowak A (2020) Homophily as a process generating social networks: insights from social distance attachment model. *J Artif Soc Soc Simul* 23(2)
- Tóth G, Wachs J, Di Clemente R, Jakobi A, Ságvári B, Kertész J (2021) Inequality is rising where social network segregation interacts with urban topology. *Nat Commun* 12(1):1143
- Uphoff N, Som A, Cbnrm B (1998) Community-based natural resource management: connecting micro and macro processes, and people with their environments
- Van Nieuwerburgh S, Veldkamp L (2010) Information acquisition and under-diversification. *Rev Econ Stud* 77(2):779–805
- Xie Y, Zhou X (2014) Income inequality in today's China. *Proc Natl Acad Sci USA* 111(19):6928–6933
- Yesuf M, Bluffstone RA (2009) Poverty, risk aversion, and path dependence in low-income countries: experimental evidence from Ethiopia. *Am J Agric Econ* 91(4):1022–1037
- Yonzan N, Mahler DG, Lakner C (2023) Poverty is back to pre-covid levels globally, but not for low-income countries. Accessed on January 29, 2025

Acknowledgements

The authors acknowledge the support from the Netherlands eScience project of the Dutch NWO, under contract 27020G08, titled “Computing societal dynamics of climate change adaptation in cities”.

Author contributions

C.D. and D.R. conceived the study and designed the model. C.D. developed the code for the model. C.D. and D.R. designed the methods and conducted the analyses. C.D. and D.R. wrote the manuscript. Both authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Ethical approval

The study does not involve human participants or their data.

Informed consent

The study does not involve human participants or their data.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1057/s41599-025-06089-9>.

Correspondence and requests for materials should be addressed to Debraj Roy.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025