Social Segregation, Misperceptions, and Emergent Cyclical Choice Patterns

Daniel M. Mayerhoffer[1,*,†], Jan Schulz[2,†]

[1]Institute for Social Science Research and Data Science Centre, University of Amsterdam
[2]Economics Department, University of Bamberg

Abstract
This paper examines the puzzle of why economic inequality has not resulted in political countermeasures to mitigate it, and proposes that the reason is due to misperceptions of economic inequality caused by segregation in social networks. We model taxation and voting behavior with an exponential income distribution and a Random Geometric Graph-type model to represent homophily, which leads to people perceiving their own income rank and income to be close to the middle. We find that people base their beliefs about mean income on a compound of the true mean and their local perception in the network, and that higher homophily causes lower implemented tax rates, which explains why redistribution preferences appear decoupled from actual inequality. In a dynamic extension, we also demonstrate that a rich set of dynamic behaviours can emerge from rational updating beliefs about efficiency. Misperceptions not only decrease redistribution in a static setting, they also hinder agents from adapting and learning towards the unbiased tax rate in a dynamic sense. As policy implications, we suggest two measures to counteract this: educating people about the actual income distribution and promoting diversity to reduce homophily.

Keywords
Inequality, redistribution, perception, bias, networks

1. Introduction

The literature on political economy has long been puzzled by the fact that, in many countries worldwide, massive increases in economic inequality have yet to prompt widespread calls for redistribution (Larcinese, 2007; Bredemeier, 2014). This phenomenon raises important questions about the relationship between economic inequality and political action. Most of the literature is based on the canonical Meltzer-Richard framework that predicts redistribution to increase with pre-tax income inequality (Meltzer and Richard, 1981). However, models of this type counterfactually assume that agents are entirely rational and have perfect knowledge about the moments of the income distribution and the redistributive capacity of the state (Bredemeier, 2014). Empirically, perceptions of one's own social status in the income hierarchy and the degree of inequality in society are heavily biased, though (cf. Jachimowicz et al., 2022, for a recent survey).

Manuscript prepared for the 35th EAEPE Conference. Preliminary draft, do not circulate.

*Corresponding author.
†These authors contributed equally. The order was chosen randomly.
d.m.mayerhoffer@uva.nl (D. M. Mayerhoffer); jan.schulz-gebhard@uni-bamberg.de (J. Schulz)
0000-0001-8841-407X (D. M. Mayerhoffer); 0000-0001-7745-3997 (J. Schulz)

As a factor contributing to this puzzle, we propose misperceptions of economic inequality caused by segregation along socioeconomic lines. We augment a conventional Downsian model of voting for redistribution with a behavioural component in an agent-based network model. It simulates agents' perceptions of inequality based on their social network and analyses the relationship between social networks, inequality perceptions, individual preferences for redistribution, and their aggregation. This model by Schulz et al. (2022b) can replicate the known stylised facts regarding misperceptions regarding society-wide inequality and inequality along gendered and racialised lines (Mayerhoffer and Schulz, 2022). We find that misperceptions can explain the puzzle of missing redistribution with now *perceived* rather than *actual* inequality driving redistributive voting behaviour, in line with empirical findings, e.g. in Choi (2019). In this sense, we find theoretical support for the Marxian notion of ideology as *necessarily false consciousness* that necessarily arises from correct inference based on skewed samples rather than cognitive failures (Schulz et al., 2022a) that induces people with low incomes to vote against their objective interests (Wisman, 2023). This ideology is not destiny, though, and we also show that agents can, in some instances, adapt and collectively learn to overcome their misperceptions.

Given the recent empirical support of beliefs about taxation efficiency in shaping public opinion, we extend the model to feature endogenous belief-formation about efficiency using adaptive expectations (Ballard-Rosa et al., 2017; García-Sánchez et al., 2020). In line with the empirical evidence, we find that poorer people tend to have more favourable beliefs about taxation efficiency than the rich (Tepe et al., 2021). Our model thus endogenously generates opinion polarization, with the electorate believing in Keynesian equitable growth at one extreme and in Okun's famous "leaky bucket" at the other opinion pole. Notably, this parsimonious extension gives rise to a rich set of dynamics in the form of monotonous or oscillatory convergence and persistent limit cycles regarding chosen tax rates. As a complement to the recent results by Di Guilmi and Galanis (2021), we thus show that endogenously evolving beliefs with fixed preferences rather than preferences that themselves evolve can lead to cyclical behaviour and convergence. Even though misperceptions regarding transfer sizes and deadweight loss persist, the unbiased benchmark might nevertheless thus emerge from the collective learning of agents. Thus, the poor can escape the initial state of low redistribution by adaptive learning. However, collective voting behaviour only converges whenever initial beliefs do not deviate too far from unbiased values, highlighting the relevance of misperceptions and homophilic segregation not only in a static sense but also in a dynamic setting.

The remainder of this paper is organised as follows: We will first introduce some background to justify our modelling assumptions, followed by a model description. In the fourth section, we describe a battery of results both regarding a static and a dynamic variant of the model and test how socioeconomic segregation can cause biases in redistributive behaviour and how persistent those are. The paper concludes with policy implications and a discussion of limitations and avenues for further research.

2. Background

Pre-tax wage distributions empirically follow an exponential distribution (Drăgulescu and Yakovenko, 2001; Silva and Yakovenko, 2004; Tao et al., 2019). We use this robust stylised fact to initialise our agents' pre-tax wages. This is not only empirically sensible but also an analytically convenient assumption since it renders the income distribution with the mean fixed to unity parameter-free. Combining this exponential distribution with the network-formation process described below implies heavily segregated social networks.

Regarding the redistribution mechanism, we broadly follow the canonical Meltzer-Richard (MR) model that features endogenous labour-leisure choices with taxation on wages disincentivising work (Meltzer and Richard, 1981). This might give rise to a taxation Laffer curve for the whole range of tax rates, i.e., tax revenues being 0 at a linear tax rate of $t = 0$ and $t = 1$, while featuring at least one revenue-maximising rate. This mechanism of deteriorating tax bases is necessary to generate an interior solution and to avoid a situation where a 100 % tax rate or "slavery of the rich" (Foley, 1967) emerges.

The MR framework assumes perfectly informed voters and predicts taxation to rise with pre-tax inequality, i.e., the median-to-mean ratio. We allow for deviations from the perfect information case and misperceptions, especially regarding the societal mean income that ultimately determines the size of the transfer. Our model nests the MR result whenever agents only rely on global information rather than their ego networks. These ego networks are explicitly modelled to capture the salient stylised facts about misperceptions of income inequality (Schulz et al., 2022b; Mayerhoffer and Schulz, 2022). In particular, we assume that agents only interact with a relatively small subset of the population (which the empirical literature on so-called Dunbar's numbers (Mac Carron et al., 2016) suggests to be about five close contacts) and within groups that are homogeneous in income, as predicted by the pronounced income homophily of empirical social networks (McPherson et al., 2001). These two properties can be formalised within a Random Geometric Graph-type of model and give rise to social networks which are, much like their real-world counterparts, sparse and strongly clustered, testifying to the external validity of the network model (Schulz et al., 2022b). Belief formation in our model is partially based on the ego networks the mechanism in Schulz et al. (2022b) generates.

In the dynamic part of the model, we build on adaptive learning, where agents base their beliefs on the previously realised values and an error-correction term regarding their forecast error (Evans and Honkapohja, 2001). Belief updating thus depends indirectly on all previously realised values and is strongly path-dependent. Adaptive expectations indeed appear to be relevant expectation formation rules in other economic domains like monetary policy (Hommes et al., 2019); however, we are not aware of applications within political economy and voting behaviour. In this sense, our model's static and dynamic parts can be interpreted as being associated with different economic paradigms: The static model emphasises equilibrium outcomes, stability and optimising behaviour, especially regarding the Laffer curve. Central tenets of neoclassical economics thus feature prominently in this model part. The dynamic extension focuses on adaptation, out-of-equilibrium dynamics and path-dependency and can thus be associated with the evolutionary approach (Foster, 1997). One way to interpret our dynamic results is, therefore, as them illustrating the potency of evolutionary adaption for self-organising behaviour, i.e., agents collectively learning the tax rate under unbiased perceptions.

3. Model

This is a non-technical overview. See Mayerhoffer and Schulz (2023) for the model and a detailed description.

Population, income distribution, and tax regime We observe a population of $1,000$ agents[1] that represents individual wage-earners. Agents are identical in all respects (including the weight for network formation and sensing introduced below), except for their income. Pre-tax wage incomes are initialised from an exponential distribution, normalised to $\lambda = 1$.
To provide microfoundations for our postulated Laffer curve mechanism, we consider the constrained optimisation problem of an agent with iso-elastic utility.[2] Agents maximise their utility $U(\cdot)$ function of the following form that is identical for all i:

$$U_i(c, y; \epsilon) = c - \frac{y_i}{1 + \frac{1}{\epsilon}} \cdot \left(\frac{\tilde{y}_i}{y_i}\right)^{1+\frac{1}{\epsilon}}, \tag{1}$$

subject to a linear budget constraint with $c = (1-t)y_i + T$, with c as consumption, y_i as the pre-tax wage income, $\tilde{y}_i$ as after-tax income, ϵ as the intensity of disutility of labour and T as the size of the transfer that agent i takes to be exogenous.[3] It follows that the optimal effort and income choice of agent i is given by $\tilde{y}_i = (1-t)^\epsilon y_i$ (Bastani and Lundberg, 2017). If all agents exhibit equal ϵ, then the mean post-tax wage income also scales in the same way with the mean pre-tax income, i.e. $N^{-1}\Sigma_i^N \tilde{y}_i = N^{-1}\Sigma_i^N (1-t)^\epsilon y_i = (1-t)^\epsilon \bar{y}$, with $\bar{y}$ as the mean pre-tax income. Given this endogenous labour-leisure choice by agents, the transfer out of taxes T is no fixed proportion of the mean pre-tax income as the tax base decreases in the tax rate, with constant elasticity ϵ:

$$T(t, \bar{y}; \epsilon) = t \cdot (1-t)^\epsilon \, \bar{y} \tag{2}$$

The above equation directly implies a Laffer curve of tax revenues and consequently transfer sizes, as also shown in Figure 1. It is straightforward to calculate the revenue-maximising tax rate as $t^* = 1/(1+\epsilon)$. Note that this revenue-maximising rate approximates the optimal rate of the poorest agent(s) with wages close to zero, as those want to maximise the size of the transfer.[4] In this sense, the revenue-maximising rate approximates the rate chosen using a Rawlsian maximin criterion (Rawls, 1974).

[1] See Schulz et al. (2022b) and the model implementation for a sensitivity analysis. This shows that the results of the network formation are robust also for larger population sizes and ego-networks, and identifies the relevant range of homophily levels.

[2] Of course, other mechanisms for s tax base that is decreasing in the tax rate are possible, e.g., tax avoidance and tax evasion or bureaucratic inefficiencies. These specific microfoundations based on iso-elastic utility are chosen for clarity of exposition. Empirically, the elasticity parameter likely results from the interplay of these effects and we incorporate reputation effects of tax evasion in our network model following Di Gioacchino and Fichera (2022) that allow us to endogenise the elasticity parameter.

[3] The assumption of an exogenous transfer T implies that agent i assumes their contribution to total tax revenue to be zero which is a plausible assumption for high N.

[4] For an exponential income distribution with $\lambda = 1$, the expected income of the poorest individual is given by the first order statistic as $1/N$ and thus approaches zero rapidly (hyperbolically) in N.

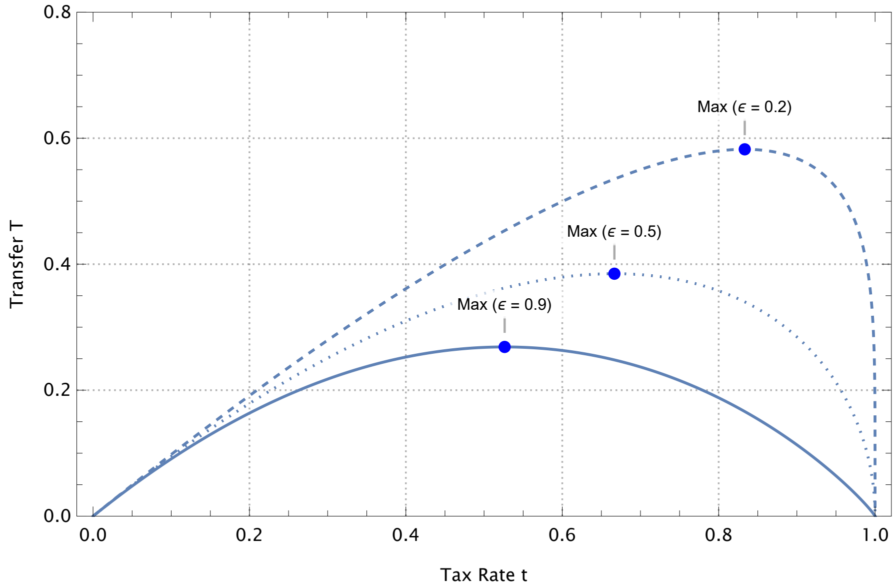

Figure 1: Laffer curves for different levels of $\epsilon \in \{0.2; 0.5; 0.9\}$, representing taxation inefficiency resulting from the disutility of work. Laffer curves feature a unique revenue-maximising rate that is further to the right the lower ϵ is.

The resulting revenue is then equally split amongst all agents, implying the following difference between pre-tax and post-tax income for each agent i:

$$\Delta y_i = -t \cdot y_i + t \cdot (1-t)^\epsilon \cdot \bar{y} \tag{3}$$

This trivially implies a threshold income $y^* = (1-t)^\epsilon \, \bar{y}$ distinguishing agents with a net benefit from the tax ($y_i < y^*$) from those with a net loss ($y_i > y^*$). y^* equals the mean pre-tax income for $\epsilon = 0$, and shrinks in ϵ and in t for $\epsilon > 0$. As a corollary on the macro-level, taxation always decreases wage inequality. To make this precise, consider the relationship between the pre- and after-tax Gini coefficient, i.e., G and $\tilde{G}$, respectively. This ratio can be expressed analytically as

$$\frac{\tilde{G}}{G} = \frac{1-t}{(1-t) + t \cdot (1-t)^\epsilon}. \tag{4}$$

It follows that $\tilde{G}/G < 1$, iff $t > 0$ and ϵ finite. The inequality-decreasing effect is mitigated for larger ϵ for any given t, since higher taxation inefficiency decreases the size of the (equalising) transfer. As a limit case, inequality is unchanged for any tax rate, if ϵ approaches infinity, that is, $\lim_{\epsilon \to \infty} \tilde{G}/G = 1$, as the transfer is then nil.

Income homophily in a Random Geometric Graph-type network To account for localised perception of tax effects, we model mutual knowledge of pre-tax incomes. Following Section 2, we assume this knowledge to predominantly exist within the agents' closest layer of interaction that is relatively homogenous in income. To account for the underlying income homophily, we employ the Random Geometric Graph-type linking procedure introduced by Schulz et al. (2022b): Each agent i draws five link-neighbours, based on weights for a potential drawee j as inversely exponentially related to the distance in pre-tax income (y) between i and j. Because knowledge about income is mutual and a relation either does or does not exist, links in the model are unweighted and undirected.

The homophily strength parameter ρ determines the extent to which proximity in income influences linkage: A value of 0 for ρ means a random network, while rising ρ means that agents with large income distances become very unlikely to be linked due to the exponential nature of the weight given by $w_{ij} = 1/\exp[\rho\,|y_j - y_i|]$. This weight is derived in Schulz et al. (2022b) from microfoundations based on a random utility framework. Schulz et al. (2022b) also show that the resulting network structure exhibits a high level of segregation, despite still being a connected graph and showing weak small-worldliness (Humphries and Gurney, 2008). Schulz et al. (2022b) prove analytically that people tend to find themselves having the median income in their ego network in line with empirical studies finding a 'middle-class bias'. To illustrate this point, we show the results for a simulated network with homophily parameter $\rho = 8$ against the empirical self-perceptions in Germany from the International Social Survey Program in Figure 2. The network simulation replicates the empirical findings remarkably well and should thus constitute a good starting point to model real-world perceptions.

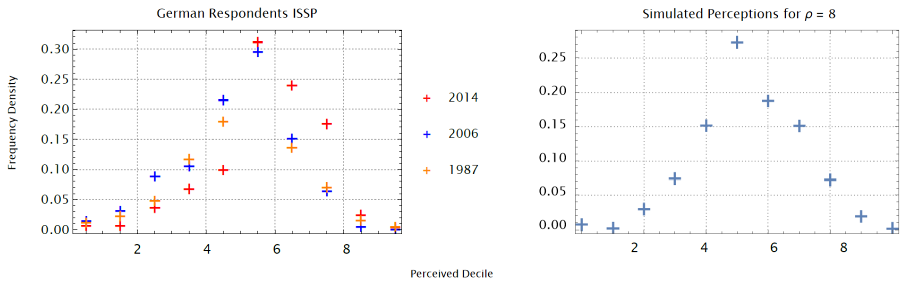

Figure 2: The figure shows self-perceptions of income deciles from a survey for German respondents in the International Social Survey Program (left panel) and as simulation outcomes in the right panel for $\rho = 8$. The middle-class bias that emerges in the simulation closely mimicks the empirical middle-class bias that empirical surveys consistently demonstrate.

Localised individual perceptions and tax rate acceptance For decisions on whether to accept any given tax rate t, we assume purely selfish motivations of agents, i.e., they accept any tax rate from which they expect a net gain for themselves. Furthermore, agents exhibit perfect, unbiased information processing and possess identical expectations about the elasticity ϵ. However, agents' sensing is imperfect. Namely, their perception of the mean income is a compound of the actual global mean $\bar{y}$ and the mean income in their ego-network l_i with a parameter $a \in [0,1]$.[5] An agent i thus believes the mean income to be $\hat{y}_i$, that is

$$\hat{y}_i = a \cdot \bar{y} + (1-a)l_i. \tag{5}$$

Consequently, the threshold income determining whether an agent expects net gains or losses from a tax rate is now individualised:

$$y_i^*(t) = (1-t)^\epsilon \cdot [a \cdot \bar{y} + (1-a) \cdot l_i] \tag{6}$$

The weight parameter $a \in [0,1]$ is identical for all agents; $a = 1$ equals perfect information, and $a = 0$ means that agents only rely on what they observe in their ego network. If $l_i < \bar{y}$, the agent does not accept some tax rates giving them a net gain; vice versa, they accept some tax rates meaning a net loss if $l_i > \bar{y}$.

Voting behaviour and aggregation As discussed above, agents accept a tax rate, whenever their income is lower than the threshold income, are indifferent if both are equal, and do not accept the tax rate if their income exceeds the threshold income. From this, we can deduce individual voting behaviour on any given tax rate t, i.e., agent i with pre-tax income y_i follows the voting rule V_i:

$$V_i = \begin{cases} 1 & \text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y} + (1-a) \cdot l_i] > y_i \\ 0 & \text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y} + (1-a) \cdot l_i] = y_i \\ -1 & \text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y} + (1-a) \cdot l_i] < y_i \end{cases} \tag{7}$$

The individual voting decision of any agent with respect to a tax rate $0 < t < 1$ is thus a function of their perception l_i, the elasticity parameter ϵ, their income y_i, the true global average income $\bar{y}$ and the weight parameter $0 \leq a \leq 1$. For $a = 1$, our model nests the standard MR framework without misperceptions. It follows that an individual tax rate t has a majority, if $V > 0$:

$$V(t, \epsilon, w, \bar{y}, \vec{l}, \vec{y}) = \sum_{i=1}^{N} V_i = \sum_{i=1}^{N} \text{sign}[(1-t)^\epsilon \cdot [a \cdot \bar{y} + (1-a) \cdot l_i] - y_i]$$

To determine the implemented tax rate t^*, we simulate the highest tax rate that a majority of agents would accept. Since the aggregator function $V(\cdot)$ monotonously decreases in t, one can simply achieve this by starting at $t = 1$ and decrease the tax rate in increments of 1 percentage point until a majority is reached.

[5] In Appendix A, we demonstrate that local perceptions can be additively decomposed into the sum of the true mean income and a bias term, which implies that network perceptions of the mean income are uniquely biased by the homophilic segregation mechanism and not by other factors such as small-sample bias. The compound rule scales the additive bias of perception downwards by a factor $(1-a)$, i.e., individual perceptions approach the true mean income $\bar{y}$ linearly in a.

Dynamic updating and beliefs about efficiency So far, we have only considered static redistribution and voting behaviour for one period. However, it is perhaps not reasonable to assume that misperceptions fully persist, even though realized transfers should conclusively show agents that their beliefs are clearly wrong. To account for this, we introduce dynamic updating via adaptive expectations for agents. This dynamic extension is motivated by the empirical finding that efficiency preferences strongly mediate the relationship between income position and redistribution (Tepe et al., 2021). The extension then allows to examine which types of initial states lead to *learnable* unbiased tax rates, i.e. collective decisions that converge to the MR benchmark with $a = 1$. We assume that agents correctly observe the transfer T and their own income y_i but sustain their misperception about the average societal income. This implies that they can only make sense of the size of the transfer by adapting their beliefs about tax efficiency ϵ. Agents first solve for ϵ from the (universally known) expression of the transfer $T = t \cdot (1-t)^{\epsilon} \cdot \hat{y}_i - t \cdot y_i$ which yields $\epsilon = \log\left(\frac{ty_i + T}{t\hat{y}_i}\right) / \log(1-t)$. We impose standard adaptive expectations for the updating process for expectations about ϵ, denoted $\epsilon^e_{i,\tau}$ at time $\tau = 0, 1, \ldots$, which reads

$$\epsilon^e_{i,\tau} = \epsilon^e_{i,\tau-1} + \lambda(\epsilon_{i,\tau-1} - \epsilon^e_{i,\tau-1}), \tag{8}$$

that is, current expectations reflect past expectations plus an error-adjustment term with factor $\lambda \in [0, 1]$, where $\epsilon_{i,\tau-1}$ is the perceived realized ϵ for agent i of the previous period (Evans and Honkapohja, 2001). Rewriting this equation and substituting the expression for the perceived realization of ϵ_i into eq. (8) yields the updating rule below

$$\epsilon^e_{i,\tau} = \lambda \cdot \frac{\log\left(\frac{t_{\tau-1} y_i + T_{t-1}}{t_{\tau-1}(a \cdot \bar{y} + (1-a)l_i)}\right)}{\log(1 - t_{\tau-1})} + (1-\lambda) \cdot \epsilon^e_{i,\tau-1}. \tag{9}$$

For $\lambda = 1$, we recover what is known in the literature as "naive expectations", i.e., expectations about ϵ are simply the perceived realized values of the previous period (Evans and Honkapohja, 2001).

4. Results

Simulation mode and validation To ensure internal validity, we simulate 100 Monte Carlo runs for each parameter combination and report aggregates of or distributions across them. The homophilic linkage mechanism as the central force in our model is also externally valid since it can reproduce the stylised empirical perception patterns of income inequality in general (Schulz et al., 2022b) and of wage-gaps (Mayerhoffer and Schulz, 2022). Furthermore, it also features a middle-class bias (Schulz et al., 2022b) because agents tend to occupy the median income rank in their ego network, as shown above. We will utilise this latter fact below. The first results concern the static setting of the model, i.e., agents not updating their beliefs, and we go on to present the results with belief adaption in a dynamic setting in the second part of the results section.

Acceptance of tax rates depending on the weight of global signal Empirically, perceived levels of inequality, not so much actual levels, drive redistributive behaviour (Choi, 2019). We examine this issue by simulating various a that represent the quality of agents' information regarding transfer size. Figure 3 shows the highest accepted tax rate given the level of tax inefficiency and weight attributed to the globally correct mean income, as opposed to their localised perception. The plot line for $a = 1$ is a benchmark for perfect sensing and an indicator of the highest tax rate that yields a net benefit for a majority of agents. However, many agents do not perceive their benefit from redistribution since the mean income in their ego networks is lower than the global mean. Thus, the accepted tax rate decreases the higher the weight put on the localised perception (i.e., the lower a). [6] We can thus replicate the finding from the literature that misperceptions are likely an important cause for the apparent lack of redistribution.

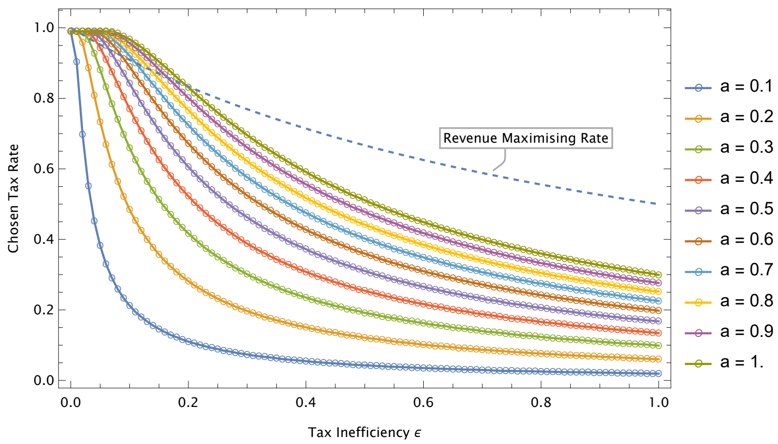

Figure 3: The figure shows the implemented tax rates for different weights for the global signal $a \in [0; 1]$ and for varying the elasticity of taxable income $\epsilon \in [0, 1]$. The simulation holds the homophily level constant at $\rho = 8$. Increasing the weight of the global signal and improving the accuracy of perceptions unanimously increases (implemented) redistribution since agents then expect higher transfers on average.

Homophily and segregation as bias-increasing factors The underestimation of one's own benefit from redistribution becomes more pronounced the higher the homophily level ρ, as Figure 4 highlights. As a corollary of the linkage procedure and the underlying income distribution, agents tend to occupy the median income in their ego network, and the median tends to be lower than the mean. However, the higher ρ, and consequently the segregation of ego networks, the smaller this difference between local mean and median grows - and the more severely agents underestimate their personal gain from redistribution. There is almost no change in redistribution preference for extreme levels of segregation ($\rho > 8$) because, in this range, the variation in incomes within an ego network shrinks, but the mean remains nearly

[6] That for low inefficiency levels, the highest tax rate capable of winning a majority exceeds the revenue maximising one is a feature of the exponential income distribution: The median income is considerably lower than the mean income; hence, a majority of agents still benefit from (almost) full redistribution even if there is some inefficiency.

unaffected. Our model thus replicates the finding that segregation tends to lower (implemented) redistribution (Franko and Livingston, 2022) and explicates the channel, i.e., segregation leads the poorer part of the population to underestimate the size of a potential transfer and benefits from redistribution.

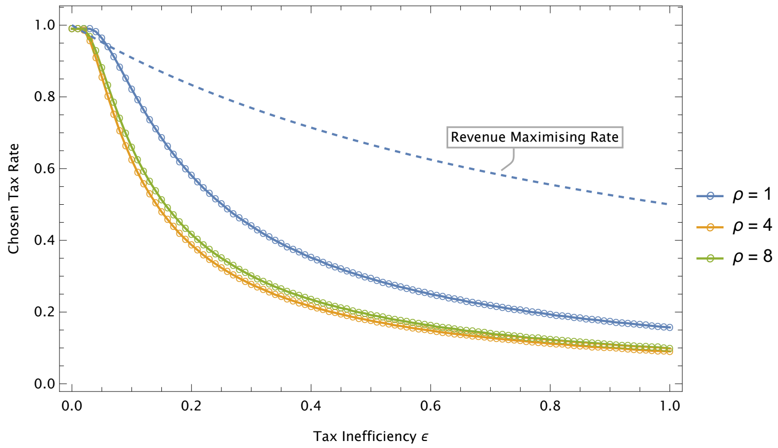

Figure 4: The figure shows the implemented tax rates for varying the elasticity of taxable income $\epsilon \in [0,1]$. For the simulation, the weight of the global signals is held constant at $a = 0.3$, and the homophily strength varies in the range $\rho \in \{1; 4; 8\}$. Increasing network segregation decreases (implemented) redistribution.

Dynamic updating and cyclical voting In a static setting with fixed beliefs about tax efficiency, biased beliefs about the mean income always lead to a downward bias in implemented redistribution. In a dynamic setting, this is not necessarily the case anymore. We let agents adapt their beliefs about taxation efficiency over time according to the updating rule in eq. (9). Figure 5 shows three exemplary belief schedules ϵ against pre-tax incomes y_i for $\rho = 8$, $a = 0.5$, $\lambda = 0.25$ and a true $\epsilon = 0.5$. Beliefs are strongly polarized, with the richest believing the redistributive system to be strongly inefficient ($\epsilon >> 1$) and the poorest even being convinced that redistribution increases the income to be redistributed ($\epsilon < 0$). Polarization increases over time until both the collectively chosen tax rates and beliefs settle in a strongly polarized state (cf. the upper left panel of Figure 6). The mechanism for this opinion polarization is rather straightforward: The homophilic network linkage leads poor agents to underestimate the mean pre-tax income and, thus, the size of the transfer, while rich agents overestimate it. Thus, the realised transfer will leave individuals puzzled, and they will ascribe the difference to an ϵ that is lower than expected for poor agents or higher than expected for rich agents. Adaptive expectations lead to convergence to a steady state whereby individual perceptions of ϵ grow more polarised: Poorer individuals assume taxation to be more efficient and become relatively more likely to accept a given tax rate, whereas richer individuals assume lower tax efficiency and become less likely to accept a given tax rate. However, due to the effects of localised perception, the poorer individuals were likely not to accept tax rates that would grant them a net benefit, whereas the richer were likely to accept rates that meant a net loss to them.

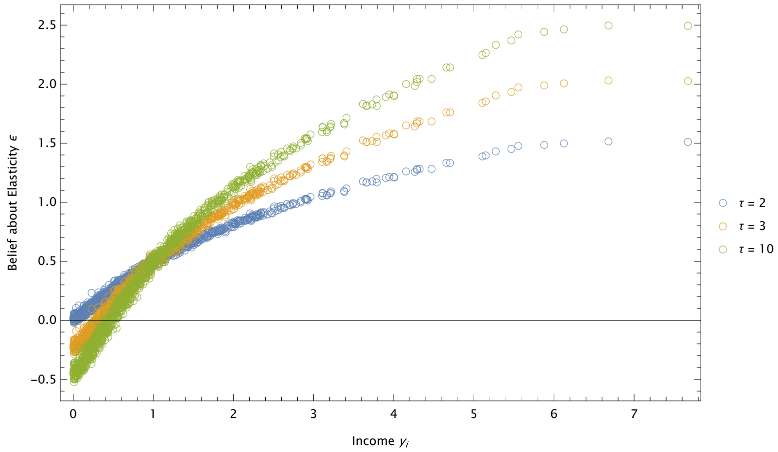

Figure 5: The figure shows agents' beliefs about the elasticity of taxable income ϵ against their pre-tax incomes y_i for $\rho = 8$, $a = 0.5$, $\lambda = 0.25$ and the true $\epsilon = 0.5$. Beliefs are polarized, with poorer agents exhibiting higher trust in tax efficiency. Opinion polarization grows, as is immediately visible by the belief schedules growing steeper through time.

Consequently, the polarization in perceptions of tax efficiency leads to a growing accord about acceptable tax rates; hence, the polarising views on efficiency actually drive the convergence to a steady state in implemented tax level. Note that in our model, in the steady state (that is already reached for $\tau = 10$ in Figure 5), almost all agents hold wrong beliefs both about the societal mean income and about tax inefficiency, that is, agents appear to learn the wrong notions adaptively. Perhaps counterintuitively, as we discuss in our elaboration on chosen tax rates below, these persistently biased beliefs do not necessarily lead to biased outcomes, in contrast to claims in the behavioural literature on the topic (cf. e.g. Caplan, 2000, on so-called "rational irrationality").

The discussed perceived elasticities give rise to three qualitative types of emergent voting behaviour: Monotonous convergence, oscillatory convergence and limit cycles. These three ideal cases are illustrated for a particular parameter constellation in Figure 6, i.e. for $a = 0.5$, $\rho = 8$ and a true $\epsilon = 0.5$ with several values for the error-correction parameter $\lambda \in \{0.25; 0.5; 0.75; 1\}$.

Monotonous convergence is illustrated in Figure 6 by the upper two panels for $\lambda = 0.25$ and $\lambda = 0.5$. For these two parameter values, there exists a self-reinforcing process leading to this convergence: As a result of a higher-than-predicted transfer following their initial vote for the static biased perception rate, poor agents adjust their beliefs about taxation efficiency ϵ downwards, leading them to vote for higher tax rates t. Since beliefs adjust only partially due to a comparatively low λ, this self-reinforcing process of higher-than-predicted transfers, a downward adjustment of beliefs about ϵ and accepting higher tax rates continues until the unbiased static rate (corresponding to $a = 1$,) is reached. In this case, transfers are not higher than expected anymore and beliefs about ϵ and the chosen t stay constant. For low λ, the unbiased static rate is thus learnable by adaptive expectations.[7]

[7]The tax rate still exhibits tiny fluctuations of around one percentage point. This artefact results from the voting

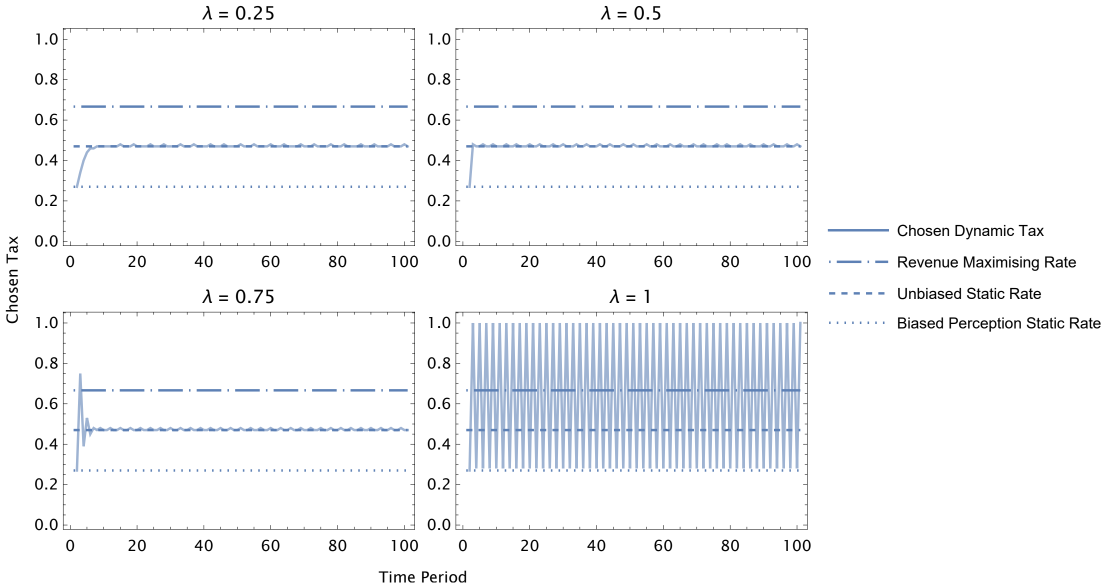

Figure 6: The figure shows the chosen tax rates for the dynamical updating process with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$. All simulations are conducted for $a = 0.5$, $\rho = 8$ and a true $\epsilon = 0.5$. The upper two panels show monotonous convergence, and the lower left panel shows oscillatory convergence to the unbiased static rate of $a = 1$. In contrast, the lower right panel shows a 2-cycle of tax rates that constantly alternate between unity and the static biased perception rate.

The same basic mechanism is behind oscillatory convergence, where chosen tax rate oscillates around the unbiased static rate with decreasing amplitude and converges eventually: Here, poor agents are initially also surprised by a transfer that is higher than predicted and adjust their beliefs accordingly. Nevertheless, since λ is higher, the adjustment is stronger and leads them to collectively vote for a tax rate that is above the unbiased static rate. For this new tax rate, the generated actual inefficiency is rather high due to the high chosen tax rate, leading poor agents to experience lower-than-expected transfer sizes and adjust their beliefs about the tax efficiency parameter ϵ upwards and vote for a lower tax rate. However, this lower tax rate is higher than the biased perception tax rate since voting also depends on agents' previous optimistic beliefs. This process of under- and overshooting continues with decreasing amplitude until the chosen rate eventually reaches the unbiased static rate and stays at this level. In both cases, the unbiased rate thus acts as a long-run attractor. We note further that the revenue maximising rate approximating the rate chosen under a Rawlsian maximin criterion does not emerge in our static or dynamical setups discussed so far. Moreover, even if it does so, it does only accidentally for very particular parameter constellations and not because of some kind of general convergence mechanism. This might imply that while the unbiased static rate is learnable even under misperceptions, the revenue maximising rate is not. To let agents self-organize around the Rawlsian maximin rule thus necessitates a different voting regime.

procedure that lets agents only vote for tax rates in increments of 1%.

Finally, for $\lambda = 1$, there exists a limit cycle, with chosen tax rates alternating between the maximum rate of unity and the biased perception static rate. Agents again initially vote based on their biased perception of the biased perception static rate. These resulting higher-than-expected transfers for poor people lead agents to fully adjust their beliefs to strongly optimistic values of ϵ, as $\lambda = 1$ and to then collectively implement the confiscatory upper limit of t, i.e., a tax rate of (almost) unity.[8] This leads to substantial objective efficiency losses, leading agents to revise their efficiency estimates and vote for a lower tax rate with, therefore, lower efficiency losses, in turn letting them revise their beliefs again in a never-ending cycle in waves of optimism and pessimism about taxation efficiency.
As the results for this parameter constellation show, the model can generate a rich set of dynamics that depend on the error-adjustment parameter λ. Perhaps unexpectedly, a lower λ or intensity of error correction tends to foster the learning of and convergence to the unbiased tax rate since low error correction levels prevent over- and undershooting around the unbiased tax rate. In other cases, oscillatory convergence or a limit cycle without any convergence emerges, providing a minimal model of electoral cycles purely based on rational belief updating (Norpoth, 2014). Electoral cycles are thus not necessarily the result of irrational bandwagon effects or an unfounded taste for change in the electorate; they can also emerge from rational learning. However, the collective voting behaviour only converges even for low values of λ (either oscillatory or monotonously) whenever initial beliefs do not deviate too far from the unbiased benchmark, as we show in Appendix B. Importantly, homophilic segregation thus not only causes misperceptions in a static setting and hinders redistribution, but it also can prevent agents from dynamic learning and self-organising around the unbiased rate.

5. Discussion

The study presented a model that suggests that misperceptions of income and, consequently, paradoxical voting behaviour can result from correct belief formation, given segregation in social networks. The model is based on a parsimonious representation of taxation and voting behaviour, and it uses a Random Geometric Graph-type model to represent the observed homophily. The results of the model support the hypothesis that individual perceptions of inequality are shaped by the people they interact with regularly and that this leads to a bias in their perceptions of how unequal the society is as a whole. We also show the implications of these misperceptions on the relationship between taxation and revenue.
The model results support our hypothesis that misperceptions of economic inequality are caused by segregation in social networks. The high level of segregation in the network structure, coupled with the agents' bias towards assuming their own income rank and income to be close to the middle, leads to a significant underestimation of the actual level of inequality in society. This has important implications for the relationship between taxation and revenue, as optimal tax rates and revenue depend on the agents' perceptions of the overall mean income. Both educational interventions and desegregationist measures might improve perceptions and thus foster redistributive measures.

[8] For computational reasons, we set an upper limit of $t = 0.9999$. The results are equivalent to higher upper bounds.

Extending our model to a dynamic setting by allowing for the endogenous adjustment of beliefs regarding taxation efficiency leads to a rich set of voting dynamics. We show that in line with some empirical evidence, the beliefs of poor people regarding taxation efficiency tend to be overly optimistic, while the rich are too pessimistic in their assessment (Berens and von Schiller, 2017). Suppose error-correction is slow and/or network segregation, and perceptions of the mean societal income are close to the true mean. In that case, voting behaviour converges to the unbiased tax rate for correct initial perceptions. Misperceptions can thus be overcome under certain conditions, and educational interventions might be essential to enable dynamic convergence and learning. For rapid error correction, cyclical voting behaviour emerges due to endogenous waves of optimism and pessimism regarding taxation efficiency that lead agents to over- and undershoot the unbiased rate. In this sense, our model might also explain empirical electoral cycles from endogenous changes in the overall degree of optimism.

Much like in our model, (perceived) self-interest is a crucial determinant of real-world voting behaviour in redistributive settings, yet, this effect is mediated by efficiency concerns that are counteracting immediate self-interest (Tepe et al., 2021). Empirically, the richer strata of the population appear to justify their opposition to redistribution by appealing to inefficiencies that might create (Berens and von Schiller, 2017). To put this finding in the context of the debate on the nexus between growth and distribution, our model lets agents converge to two distinct positions: The richer part of the population appears to follow a neoclassical argument of the "leaky bucket" (Okun, 1975) of redistribution, in the sense that taxation on labour disincentivizes work and lets the taxable income that is to be distributed shrink. This is also the argument that motivated our microfoundation for the ϵ parameter in the form of a labour-leisure trade-off. By contrast, poorer people appear to converge to a Keynesian position that is untenable in a static sense but might be sensible in the present, dynamic setting: Redistribution from the richer to the poorer strata of the population increases aggregate consumption, hence strengthening effective demand and contributes to growth, since the poor exhibit much larger marginal propensities to consume (Pressman, 1997).[9] The long-standing nature of this debate without clear consensus renders it plausible that people exhibit consistently polarized opinions on the matter. Our model robustly predicts this polarization for different parameter constellations, which presents an opportunity for empirical validation. Yet, apart from case studies like Berens and von Schiller (2017) that are limited in scope but broadly confirm our theoretical predictions, we are not aware of any attempts to test examine the relationship between income position and beliefs about efficiency empirically.

Even though our results rely on rather specific assumptions regarding the Laffer curve, income distribution, perception and expectation formation and aggregation mechanism, they provide a robust joint explanation of why perceptions are so skewed and how they might mediate the nexus between inequality and redistribution. A model extension could include the role of media, education, or political polarisation. Furthermore, the model can also be adapted to other voting and tax regimes with progressivity to examine the robustness of the results. While our perceptions based on homophilic networks appear externally valid, belief formation also lacks a

[9]Empirically, the literature on the inequality-growth nexus has been somewhat ambiguous and dependent on the chosen sample and time-period (Flechtner and Gräbner, 2019). For a literature review of the impact of inequality on consumption and the relevance of consumption emulation mediating this effect, see Schulz and Mayerhoffer (2021).

purposive element, e.g. regarding communication in networks. With our model, we hope to provide a first benchmark that can be extended along those lines.

References

Ballard-Rosa, C., Martin, L., and Scheve, K. (2017). The structure of American income tax policy preferences. *Journal of Politics*, 79(1):1–16.

Bastani, S. and Lundberg, J. (2017). Political preferences for redistribution in Sweden. *Journal of Economic Inequality*, 15:345–367.

Berens, S. and von Schiller, A. (2017). Taxing higher incomes: what makes the high-income earners consent to more progressive taxation in Latin America? *Political Behavior*, 39:703–729.

Bredemeier, C. (2014). Imperfect information and the Meltzer-Richard hypothesis. *Public Choice*, 159(3):561–576.

Caplan, B. (2000). Rational irrationality: a framework for the neoclassical-behavioral debate. *Eastern Economic Journal*, 26(2):191–211.

Choi, G. (2019). Revisiting the redistribution hypothesis with perceived inequality and redistributive preferences. *European Journal of Political Economy*, 58:220–244.

Di Gioacchino, D. and Fichera, D. (2022). Tax evasion and social reputation: The role of influencers in a social network. *Metroeconomica*, 73(4):1048–1069.

Di Guilmi, C. and Galanis, G. (2021). Convergence and divergence in dynamic voting with inequality. *Journal of Economic Behavior & Organization*, 187:137–158.

Drăgulescu, A. and Yakovenko, V. M. (2001). Evidence for the exponential distribution of income in the USA. *The European Physical Journal B*, 20(4):585–589.

Evans, G. and Honkapohja, S. (2001). Expectations, economics of. In Smelser, N. J. and Baltes, P. B., editors, *International Encyclopedia of the Social & Behavioral Sciences*, pages 5060–5067. Pergamon, Oxford.

Flechtner, S. and Gräbner, C. (2019). The heterogeneous relationship between income and inequality: a panel co-integration approach. *Economics Bulletin*, 39(4):2540–2549.

Foley, D. K. (1967). Resource allocation and the public sector. *Yale Economic Essays*, 7:45 – 98.

Foster, J. (1997). The analytical foundations of evolutionary economics: From biological analogy to economic self-organization. *Structural Change and Economic Dynamics*, 8(4):427–451.

Franko, W. W. and Livingston, A. C. (2022). Economic segregation and public support for redistribution. *Social Science Journal*, 59(3):375–393.

García-Sánchez, E., Osborne, D., Willis, G. B., and Rodríguez-Bailón, R. (2020). Attitudes towards redistribution and the interplay between perceptions and beliefs about inequality. *British Journal of Social Psychology*, 59(1):111–136.

Hommes, C., Massaro, D., and Salle, I. (2019). Monetary and fiscal policy design at the zero lower bound: Evidence from the lab. *Economic Inquiry*, 57(2):1120–1140.

Humphries, M. D. and Gurney, K. (2008). Network 'small-world-ness': a quantitative method for determining canonical network equivalence. *PloS One*, 3(4):e0002051.

Jachimowicz, J. M., Davidai, S., Goya-Tocchetto, D., Szaszi, B., Day, M. V., Tepper, S. J., Phillips, L. T., Mirza, M. U., Ordabayeva, N., and Hauser, O. P. (2022). Inequality in researchers' minds: Four guiding questions for studying subjective perceptions of economic inequality. *Journal of Economic Surveys*, in print.

Larcinese, V. (2007). Voting over redistribution and the size of the welfare state: the role of turnout. *Political Studies*, 55(3):568–585.

Mac Carron, P., Kaski, K., and Dunbar, R. (2016). Calling Dunbar's numbers. *Social Networks*, 47:151–155.

Mayerhoffer, D. M. and Schulz, J. (2022). Perception and privilege. *Applied Network Science*, 7(1):1–25.

Mayerhoffer, D. M. and Schulz, J. (2023). Inequality-Perception: https://github.com/mayerhoffer/Inequality-Perception.

McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444.

Meltzer, A. H. and Richard, S. F. (1981). A rational theory of the size of government. *Journal of Political Economy*, 89(5):914–927.

Norpoth, H. (2014). The electoral cycle. *PS: Political Science & Politics*, 47(2):332–335.

Okun, A. M. (1975). *Equality and Efficiency: The Big Trade–off*. WW Norton and Co.

Olley, G. S. and Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297.

Pressman, S. (1997). Consumption, income distribution and taxation: Keynes' fiscal policy. *Journal of Income Distribution*, 7(1):29–44.

Rawls, J. (1974). Some reasons for the maximin criterion. *American Economic Review*, 64(2):141–146.

Schulz, J., Mayerhoffer, D., and Gebhard, A. (2022a). Soziale Vergleiche und wahrgenommene Ungleichheit. *Der öffentliche Sektor: The Public Sector*, 48(2):15–18.

Schulz, J. and Mayerhoffer, D. M. (2021). A network approach to consumption. *BERG Working Paper Series*, 173.

Schulz, J., Mayerhoffer, D. M., and Gebhard, A. (2022b). A network-based explanation of inequality perceptions. *Social Networks*, 70:306–324.

Silva, A. C. and Yakovenko, V. M. (2004). Temporal evolution of the "thermal" and "superthermal" income classes in the USA during 1983–2001. *EPL (Europhysics Letters)*, 69(2):304.

Tao, Y., Wu, X., Zhou, T., Yan, W., Huang, Y., Yu, H., Mondal, B., and Yakovenko, V. M. (2019). Exponential structure of income inequality: evidence from 67 countries. *Journal of Economic Interaction and Coordination*, 14(2):345–376.

Tepe, M., Paetzel, F., Lorenz, J., and Lutz, M. (2021). Efficiency loss and support for income redistribution: Evidence from a laboratory experiment. *Rationality and Society*, 33(3):313–340.

Wisman, J. D. (2023). Why Ideology Exists. *Journal of Economic Issues*, 57(1):200–217.

A. Appendix: Perceived Mean Incomes

Agents form their perceptions of the mean income within their ego network based on the incomes they observe, including their own. Let the adjacency matrix of the underlying networks be given as A, with the $A_{ij} = 1$ if i observes j and 0 otherwise. Define matrix $\tilde{A} = A + \mathbb{I}_{N \times N}$, with $\mathbb{I}_{N \times N}$ being the identity matrix with dimensions $N \times N$. Entries of 1 in row i then denote the position j of incomes agent i observes. Define further $\vec{y}$ as the vector of incomes and $\vec{p}$ as the vector of perceived incomes. Finally, normalise all entries in $\tilde{A}$ such that rows sum to unity to get B. A row i then still has either 0 or $1/d_i$ entries, with d_i being one plus the degree of agent i. It now has to hold that

$$B\vec{y} = \vec{p}. \tag{10}$$

Writing this relationship as a sum, we get

$$l_i = \sum_{j=1}^{N} b_{ij} \cdot y_j. \tag{11}$$

Since all b weights sum to unity by definition, we can apply the decomposition by Olley and Pakes (1996) to write

$$l_i = \bar{y} + (N-1) \cdot cov(b_{ij}, y_j). \tag{12}$$

This implies that perceptions are additively separable into the true mean and a bias term that depends on the correlation between weights and incomes. It is easy to see that perceptions l are correct, whenever weights and incomes are uncorrelated and $cov(\cdot) = 0$, consequently. Since the sign and size of the covariance term is uniquely determined by the income position for homophilic segregation, distorted perceptions thus also arise from this homophilic network formation. Consider now the compound rule for agents' beliefs about the societal mean income, i.e., $\hat{y}_i$. The compound rule is given by

$$\hat{y}_i = a \cdot \bar{y} + (1-a)l_i. \tag{13}$$

Substituting l_i from eq. (12) yields

$$\hat{y}_i = a \cdot \bar{y} + (1-a)(\bar{y} + (N-1) \cdot cov(b_{ij}, y_j)) \tag{14}$$

and rearranging finally implies

$$\hat{y}_i = \bar{y} + (1-a) \cdot (N-1) \cdot cov(b_{ij}, y_j). \tag{15}$$

Compared to the purely localised perceptions given in eq. (12), the compound rule in eq. (12) thus scales the bias term down by $(1-a)$.

B. Appendix: Dynamic Voting Behaviour for Varying a

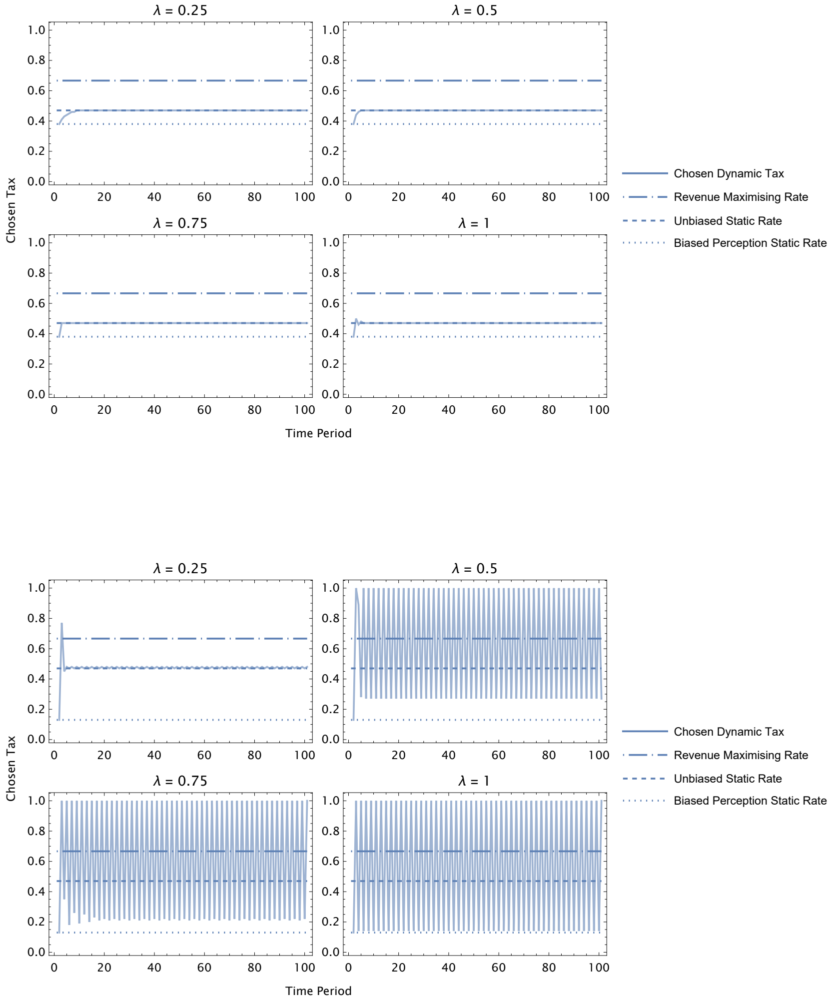

Figure 7: The figure shows the chosen tax rates for a dynamical updating process with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$. All simulations are conducted for $\rho = 8$ and a true $\epsilon = 0.5$. The upper four panels correspond to the case of "high" $a = 0.75$, while the lower four panels report the results for "low" $a = 0.25$. It is immediately obvious that strong misperceptions about the true societal mean $a < 0.5$ make convergence and learning more unlikely, as the voting behaviour for the low a is consistenly oscillating without any sign of convergence for all but one case. By contrast, almost immediate convergence is achieved in all cases for high a. Improving the accuracy of beliefs a can thus foster convergence to the unbiased state.