Public Choice (2026) 207:611–644
https://doi.org/10.1007/s11127-025-01349-3

Social segregation, misperceptions, and emergent cyclical choice patterns

Daniel M. Mayerhoffer[1] · **Jan Schulz**[2]

Received: 18 July 2023 / Accepted: 13 October 2025 / Published online: 27 December 2025

Abstract

We examine the puzzle of why economic inequality has not resulted in political measures to mitigate it, and we propose that the reason is that segregation in social networks causes people to misperceive inequality. We model taxation and voting behavior with an exponential income distribution and a Random Geometric Graph–type model to represent homophily, which leads people to perceive their own income rank and income to be close to the middle. We assume 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 behaviors can emerge from rational updating beliefs about efficiency and societal average income. The modeled misperceptions not only decrease redistribution in a static setting but hinder agents from adapting and learning the unbiased tax rate in a dynamic setting.

Keywords Inequality · Redistribution · Perception · Segregation · Bias · Networks

1 Introduction

Scholars of political economy have long been puzzled that in many countries worldwide, massive increases in economic inequality have yet to prompt widespread calls for redistribution (Larcinese 2007; Bredemeier 2014). This puzzling fact raises important questions about the relationship between economic inequality and political action. Most of the literature is

Daniel M. Mayerhoffer and Jan Schulz contributed equally.

✉ Daniel M. Mayerhoffer
d.m.mayerhoffer@uva.nl

Jan Schulz
jan.schulz-gebhard@uni-bamberg.de

1 Institute for Social Science Research, University of Amsterdam, Amsterdam, The Netherlands

2 Economics Institute, University of Bamberg, Bamberg, Germany

based on the canonical Meltzer-Richard (MR) framework, which predicts that redistribution increases 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 (see Jachimowicz et al. 2022, for a recent survey).

We propose that part of the explanation for the puzzle is that segregation along socioeconomic lines causes misperceptions of economic inequality (Georgiadis and Manning 2012). A growing literature develops models to show that these misperceptions might explain why the electorate does not vote for increased redistribution to mitigate the increasing pre-tax inequality (Bredemeier 2014; Schokkaert and Truyts 2017; Iacono and Ranaldi 2021). However, these models are often both very stylized and static—and they do not explain where the misperceptions come from in the first place.

To close this gap, we augment a conventional Downsian model of voting for redistribution with a behavioral component in a network model featuring agents with income and belief heterogeneity. It simulates agents' perceptions of inequality based on their social network and analyzes the relationship between social networks, inequality perceptions, individual preferences for redistribution, and their aggregation. The underlying network model by Schulz et al. (2022b) can replicate the known stylized facts regarding misperceptions of societal inequality and inequality along gendered and racialized lines (Mayerhoffer and Schulz 2022).

We find that misperceptions help to explain[1] the puzzle of relatively low redistribution, with *perceived* rather than *actual* inequality driving redistributive voting behavior, in line with empirical findings (for example, in Choi 2019).

Our model implies that it is not cognitive failures but the structure of human interactions mediated by homophily that lets agents vote against their objective material interests (Schulz et al. 2022a). Thus, we propose that "echo chambers" and apparently paradoxical voting behavior are directly linked and contribute to the growing literature on the implications of (lack of) epistemic diversity (Guerrero 2021). In a dynamic extension, we also show that these misperceptions are not destiny, and we demonstrate that agents can sometimes adapt and collectively learn to overcome their misperceptions.

Given the recent empirical evidence that beliefs about taxation efficiency shape redistributive preferences (Ballard-Rosa et al. 2017; Tepe et al. 2021; Barnes 2022) and that those beliefs change rapidly (Georgiadis and Manning 2012), we extend our model to feature endogenous belief formation about efficiency and average societal income as two ideal types of adaptive expectations. We find that poorer people tend to have more favorable beliefs about taxation efficiency than the rich. This is because the rich initially expect societal mean income and, thus, the transfer size to be much higher than it actually is. The only way to make sense of this is to assume that tax enforcement is very costly and leads to substantial deadweight loss. The reverse is true for poor agents, and they revise their estimates of taxation inefficiency downward. Our model thus endogenously generates opinion polarization, with voters at one pole believing in efficient redistribution while people at the other one agree with Okun's famous assumption of a "leaky bucket."

[1] This paper does not attempt a how-actually explanation but a how-possibly explanation. See Sections 4 and 5 for details.

Interpreted in this way, our model thus outlines a novel channel by which opinion polarization along socioeconomic lines—which has been consistently documented empirically (see Arndt 2020, forareview)—can be explained. Our model suggests that the rich tend to be more economically conservative than the poor not because of self-interest but because they extrapolate correctly from biased samples. 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 recent theoretical 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 behavior and convergence.

Similarly, if we let agents adapt their beliefs about average societal income, this unanimously leads to convergence to the unbiased state and more redistribution. Therefore, the poor can escape the initial state of low redistribution by learning adaptively either about taxation efficiency or about the taxable societal income to be redistributed. However, collective voting behavior might only converge (or at least might converge more rapidly) when initial beliefs do not deviate too far from unbiased values, highlighting the relevance of misperceptions and homophilic segregation not only in a static setting but also in a dynamic setting.

The remainder of this paper is organized as follows: We first provide some background to justify our modeling assumptions. Second, we describe the model. In the fourth section, we describe a battery of results regarding the model's static and dynamic variants and investigate how socioeconomic segregation can cause biases in redistributive behavior and how persistent those are. Finally, the paper discusses the implications and limitations of these findings and suggests avenues for further research.

2 Background

Our model attempts to synthesize stylized facts concerning persistently biased beliefs about inequality, regularities in networks and distributions, and voting behavior in a parsimonious model. We then situate our model components within the literature.

2.1 Income distribution

Since we are interested in voting behavior resulting from inequality in the pre-tax distribution, we take the pre-tax wage distribution to be exogenous. Empirically, pre-tax wage distributions are well approximated by an exponential distribution (Drăgulescu and Yakovenko 2001; Silva and Yakovenko 2004; Tao et al. 2019). We use this robust stylized fact to initialize our agents' pre-tax wages. This assumption is empirically sensible and analytically convenient since it renders the income distribution with the mean set to unity parameter-free. Combining this exponential distribution with the network-formation process described below implies heavily segregated social networks. However, the exponential distribution has a fixed level of inequality, with a Gini coefficient of $G \approx 0.5$. To examine the effect of different levels of inequality on redistribution, we use a lognormal distribution and vary its

σ parameter in $\sigma \in [0.01, 1]$ to see how homophilic segregation and misperceptions might mediate the link between changes in inequality and redistribution.[2]

2.2 Politics of inequality

For the redistribution mechanism, we broadly follow the canonical MR model (Meltzer and Richard 1981). In the MR framework, implemented tax rates follow the median voter's choice of a proportional tax on all incomes that is then fully transferred to all agents as a lump sum. Intuitively, mean pre-tax income determines the size of the transfer, while median income determines the choice of the median voter. Therefore, the major prediction of the MR model is that the demand for redistribution rises in pre-tax inequality proxied by the mean-to-median ratio. The empirical literature testing this hypothesis has been inconclusive and ambiguous. While Borge and Rattsø (2004), Meltzer and Richard (1983), and Milanovic (2000) find a positive link between pre-tax inequality and redistribution, Georgiadis and Manning (2012) find no significant relationship, and Lindert (1996) and Scervini (2012) even detect a negative relationship.

To make sense of these contradictory results, several scholars have extended the MR framework. First, as Larcinese (2007) demonstrates, turnout rates increasing in pre-tax income might shift the identity of the median voter and thus decrease redistribution compared to that in the MR framework, which assumes full voter turnout. Second, within the MR model, implemented redistribution is fully determined by the voting decision of the median voter, abstracting from the supply side of politics. But as Bartels (2009) argues, political leaders might be influenced by the lobbying of the financially powerful and disregard the preferences of the majority. Finally, there are models of other-regarding (or altruistic) preferences deviating from the MR assumption of purely self-interested voters (see Dimick et al. 2018, for a review). As a prominent example, Alesina and Angeletos (2005) augment the baseline MR model with fairness considerations and show that rising inequality decreases implemented redistribution. We abstract from these considerations to isolate the impact of segregation-induced misperceptions on redistributive preferences: All agents in our model are purely self-interested, the demand side of politics fully determines policy outcomes, and there is always full voter turnout.

2.3 Social segregation and misperceptions

To explain the dampened response of redistributive policy to increasing pre-tax inequality, we follow Bredemeier (2014) in allowing for deviations from the perfect-information case and for misperceptions, especially regarding mean societal income, which ultimately determines the size of the transfer. Our model nests the MR result whenever agents rely on global information only rather than their ego networks. The major idea is that agents form their beliefs based on social sampling—that is, they extrapolate from their sample to the population mean based on everyday interactions (Enke 2020). This mechanism is well documented in the psychological literature (Galesic et al. 2012; Sumaktoyo et al. 2022). As Sands (2017) and Sands and de Kadt (2020) show, exposure to wealth and inequality affects demand for redistribution, in line with the outlined mechanism.

[2] We thank an anonymous reviewer for their very helpful suggestion to examine the effects of varying levels of inequality.

With social sampling, the accuracy of beliefs is, therefore, a function of how representative the sample is for the overall population. Schulz and Mayerhoffer (2024) demonstrate theoretically that self-anchoring—the fact that agents always include their own income in their perception set—lets the majority of agents underestimate the global mean income and thus decreases implemented redistribution relative to the full-information case even for fully random social samples. Empirically, local conditions seemingly shape inequality perceptions, as Xu and Garand (2010) and Minkoff and Lyons (2019) show for the US. Residential segregation should render social samples unrepresentative of the aggregate economy. Franko and Livingston (2022) show that people favor less redistribution when their local surroundings are more segregated. Similarly, Windsteiger (2022) demonstrates empirically that implemented redistribution is correlated with residential segregation and the resulting misperceptions, and she provides a theoretical rationale based on a sorting model. The mechanism she describes is very close to the one we employ but does not consider the behavior and perceptions of individual agents within social networks.

2.4 Segregation in networks

In contrast to the above-mentioned studies of segregation, which look at residential segregation, we adopt a more general view of segregation by income within social networks. This segregation might be induced by residential segregation but can also emerge from, for example, shared workplaces, educational segregation, or segregation in leisure activities. The direct contacts of agents are called ego networks. Ego networks capture the salient stylized facts about misperceptions of income inequality (Schulz et al. 2022b; Mayerhoffer and Schulz 2022). In particular, we assume that agents only interact with a small subset of the population (which the empirical literature on Dunbar's number (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 empirically observed social networks (McPherson et al. 2001). These two properties can be formalized within a Random Geometric Graph–type model and give rise to social networks that, much like their real-world counterparts, are 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 generated by the mechanism in Schulz et al. (2022b). Similarly, Schokkaert and Truyts (2017) show that reference dependence, based on homophilic attachment, can shape beliefs about the relative importance of effort and circumstances for income differences and might thus shape redistributive preferences. Our model complements theirs, as we propose homophilic segregation to explain (mis)perceptions of the income that is to be redistributed, while Schokkaert and Truyts (2017) focus on the origins of this income.

2.5 Political business cycles and expectations

In the dynamic part of the model, we follow the foundational model of Nordhaus (1975) and assume adaptive learning: Agents base their beliefs on previously realized values and an error-correction term regarding their forecast error (Evans and Honkapohja 2001).[3] Since our model builds on the effects of the empirically documented, persistently biased percep-

[3] Including this dynamic voting component and considering varying levels of inequality are two major extensions of the paper compared to Mayerhoffer and Schulz (2023).

tions of the electorate, it seems appropriate to not assume what Nordhaus et al. (1989) critically term "ultrarational" voters with rational expectations. Experimentally, expectations indeed appear to deviate from the rational-expectation benchmark, with adaptive expectations being empirically relevant expectation-formation rules in economic domains such as monetary policy and financial markets (Anufriev and Hommes 2012; Hommes 2011; Hommes et al. 2019).[4] In Nordhaus (1975), a strategically acting government maximizing its chance of reelection leads to cyclical behavior of the unemployment and inflation rates. The cyclical fluctuations are driven by the electoral calendar: Before an election, opportunistic policy makers have an incentive to boost employment through deficit spending, which leads to high inflation after the election, necessitating austerity measures to bring inflation (and, consequently, employment) down. Cycles in political approval for left- and right-leaning parties have been documented extensively (Aguiar-Conraria et al. 2012, 2013; Lebo and Norpoth 2007; Merrill et al. 2008; Norpoth 1995, 2014).

With our model, we aim to answer the call by Dubois (2016) to integrate inequality and redistribution into the study of political business cycles. Our dynamic extension represents a minimal model of inequality-related political business cycles, as model dynamics indeed can feature cyclical behavior. We deviate from the extant literature in two ways, however: First, our model assumes single-issue voting and does not need to assume a strategically acting government to create cycles. Second, cycles emerge purely from endogenous waves of optimism and pessimism about the state's ability to efficiently redistribute income and thus come from the demand side, rather than the supply side, of politics.

3 Model

This is a nontechnical overview. See the repository (Mayerhoffer and Schulz 2025) for the model implementation and a detailed description.

3.1 Sequence of events

The events in the model unfold according to the following sequence, also visualized in Fig. 1:

1. All agents i are endowed with gross income y_i.
2. Agents form linkages in the network based on their gross income y_i and their perceptions of the societal mean gross income based on their ego networks and a global, correct signal.
3. Based on their perceptions and their own gross income, all agents vote on a flat tax rate t. The highest tax rate that gives a net benefit to agents is implemented. This decision takes government inefficiencies into account (see below).
4. The government collects taxes based on the chosen tax rate. Tax collection leads to inefficiencies since there might be efforts to avoid taxes and since the bureaucracy collecting taxes needs to be financed.

[4] This experimental evidence also points to persistently heterogeneous expectations. An interesting model extension in further research would thus be to introduce an evolutionary selection mechanism in expectation formation to account for this heterogeneity.

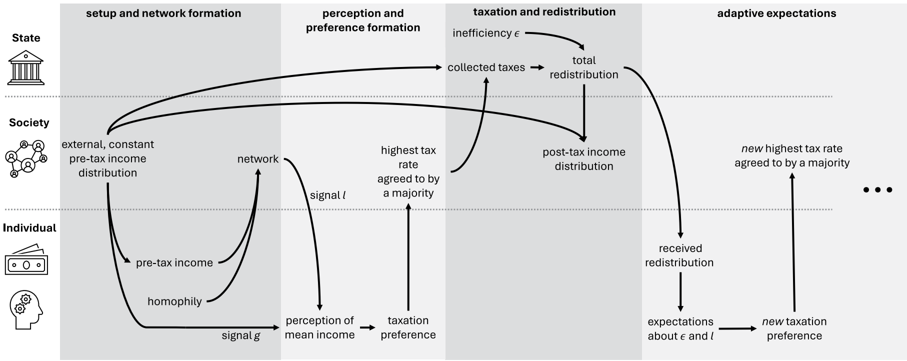

Fig. 1 Simulation flow showcasing individual choices and aggregation mechanisms

5. The net revenue is paid to agents as a lump-sum transfer of equal size. Agents receive less than they paid in taxes since part of the tax revenue is paid to finance tax collection, giving rise to a Laffer curve.
6. For the dynamic part, agents either update their beliefs about the local mean gross income or the taxation inefficiency and vote according to their beliefs.
7. Agents update their beliefs again based on the realized transfer and vote once more.

We discuss the events and assumptions in more detail in the order outlined above.

3.2 Population, income distribution, and tax regime

We observe a population of 1000 agents[5] that represent individual wage earners. Agents are identical in all respects (including the weight for network formation and sensing introduced later in this section) except their gross income. Pre-tax wage incomes are initialized from either an exponential distribution, normalized to a mean income $\lambda = 1$, or a lognormal distribution with the dispersion parameter $\sigma \in [0.01, 1]$. The former is intended to approximate the empirically relevant cases, while the latter is designed to capture the impact of inequality on segregation and voting behavior, which an exponential distribution—with a fixed Gini coefficient—cannot cover.

3.3 Income homophily in a Random Geometric Graph–type network

To account for localized perception of tax effects, we model mutual knowledge of gross incomes y. Following Sect. 2, we assume this knowledge predominantly exists within the agents' closest layer of interaction that is relatively homogeneous in income. The main idea here is that the agents' choice of links in the social network is homogeneous and that information about the moments of the income distribution—in particular, average income—is retrieved from these skewed perception sets, much like in the model of Schokkaert and Truyts (2017).

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

To represent homophily in income, we employ the network model introduced by Schulz et al. (2022b). Each agent draws five link neighbors from the set of all other agents, with the absolute weight[6] for agent i to draw agent j as a link neighbor (or vice versa) being determined by the homophily strength and the distance between i's and j's gross incomes:

$$w_{ij} = 1/\exp[\rho\ |y_j - y_i|] \tag{1}$$

The homophily-strength parameter ρ determines the extent to which proximity in income influences linkage: $\rho = 0$ denotes a random network (all weights are 1), while any $\rho > 0$ means that weights decrease exponentially in income difference—that is, the agents with the most distant incomes are very unlikely to become link neighbors. The higher ρ, the stronger this mechanism of putting a disproportionately high weight on agents with close incomes. This homophily parameter is derived in Schulz and Mayerhoffer (2023) from microfoundations based on a random-utility framework; it is also the same for all agents in a simulation run but varies between runs. Because knowledge about income is mutual and a relation either does or does not exist, links in the model are unweighted and undirected. An agent's ego network thus represents the aforementioned closest layer of interaction, which partially represents the basis for their belief formation.

Schulz et al. (2022b) identify the resulting network structure to replicate features of a Random Geometric Graph, even though the network-generation process is different. They find that it exhibits a high level of segregation, despite still being a connected graph and showing weak small-worldliness (Humphries and Gurney 2008). Figure 2 reaffirms this finding for $\rho = 8$: Networks are strongly segregated by income (with income ranks going from poor to rich) as long as inequality in the income distribution is sufficiently large. Indeed, even though all depicted networks are homophilic, the left panel in Fig. 2 for a log-normal income distribution with dispersion parameter $\sigma = 0.01$ and thus a Gini coefficient

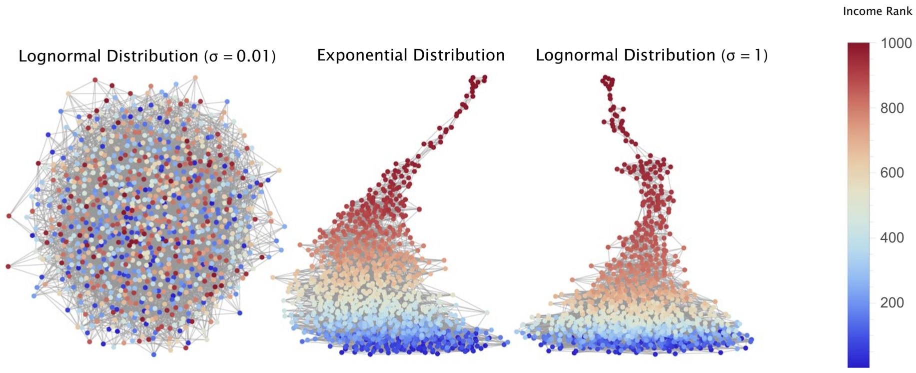

Fig. 2 The figure shows three graphs colored according to income rank (from poor to rich), with $\rho = 8$ as the homophily strength and with differing initial distributions. For low levels of income inequality (left panel), no segregation emerges, but if inequality is sufficiently high (middle and right panels), the graph is strongly segregated, with the rich being most selective in their choice of link neighbors. The graphs were plotted using Mathematica 13.0's "SpringElectricalEmbedding" routine

[6] The relative weights for j to be drawn by i and for j to be drawn by i differ since they also depend on the absolute weight of all other agents for i versus j being the one drawing the link.

of $G \approx 0.005$ mimics a fully random graph without any segregation according to income. By contrast, for an exponential income distribution (middle panel, $G \approx 0.5$) and a lognormal income distribution with $\sigma = 1$ (right panel, $G \approx 0.52$), a strongly segregated graph emerges, with the richest agents at the top being by far the most selective.

3.4 Balanced budget, bureaucratic inefficiencies, and Laffer curve

For the redistribution mechanism, we broadly follow the MR framework and integrate taxation inefficiencies into a model of proportional taxation with equal-sized lump-sum transfers. We assume agents are fully aware of these mechanisms, even though they might misperceive the parameters, particularly the mean gross income. To capture taxation inefficiencies, we integrate costs of redistribution into the model that are monotonously increasing in the tax rate and revenue.

This mechanism gives rise to a Laffer curve for the whole range of tax rates—that is, tax revenues are 0 at a linear tax rate of $t = 0$ or $t = 1$—while featuring at least one revenue-maximizing rate. This mechanism of a deteriorating tax base is necessary to generate an interior solution and to avoid a situation in which a 100% tax rate (or "slavery of the rich" (Foley 1967)) emerges. Typically, Laffer curves derive from the distortionary effects of (labor) taxation, with taxes decreasing the labor supply and, consequently, the taxable gross labor income for a given tax rate (see, for example, Bastani and Lundberg 2017). In these models, gross (taxable) incomes are thus endogenous to the model and partially determined by the tax rate. Yet since the tax rate is the major simulation *outcome* of our model and is determined by agents' perceptions, which, in turn, are based on gross incomes, we need to assume pre-tax wages to be exogenous. To generate a Laffer curve and make the results comparable to those in the MR framework, we opt for the inefficiency to work at the level of the government and we take pre-tax wages as exogenous. With this assumption, we abstract from any labor-supply responses and assume that inefficiencies arise only at the level of tax enforcement.

We assume that the government's budget is balanced—that is,

$$R(t) - \sum_i T_i(t) - C(t) = 0, \tag{2}$$

where R is government revenue, $\sum_i T$ is the sum of transfers to all agents i, and C is the costs of redistribution, which are a function of the proportional tax rate $0 \leq t \leq 1$. From the assumption of a proportional tax rate, it follows that

$$R(t) = Nt\bar{y}, \tag{3}$$

where N is the number of agents and $\bar{y}$ the mean pre-tax income.

For the cost function, we assume that all costs of redistribution arise from the necessary bureaucracy as a function of the revenue R and the tax rate t. Following the notion that tax collection becomes increasingly difficult when more tax revenue is collected, we assume that costs are proportional to revenues. We also assume that enforcement costs are an increasing function of tax resistance and avoidance efforts in the population but that all desired taxes are eventually fully enforced (that is, $R(t) = Nt\bar{y}$ indeed holds). In particular, we adopt the following functional form:

$$C(t) = E(t) \cdot R(t) \tag{4}$$

Here, the necessary effort to combat tax avoidance is E, $\partial E/\partial t \geq 0$, and $0 \leq E \leq 1$. The last of these requirements ensures that tax collection costs never exceed tax revenues; otherwise it would be doubtful that any tax would be collected in the first place. E can be interpreted as a measure of (the necessary effort to combat) tax resistance; it is zero for a zero tax rate and reaches a maximum of one for $t = 1$ since tax resistance also increases in t. The general idea here is that agents' efforts to hide their income rationally increase with the tax rate, making it more costly to identify gross taxable income.[7] The function

$$E(t) = (1 - (1 - t)^{\epsilon}) \tag{5}$$

fulfills the criteria for the enforcement cost function E. Tax resistance is zero for $t = 0$ and maximal for $t = 1$. We introduce $\epsilon > 0$ as a parameter to measure the intensity of resistance. For $\epsilon = 1$, tax resistance reduces to t, but for a different ϵ, the function becomes nonlinear. For the cost function, it follows that

$$C(t) = (1 - (1 - t)^{\epsilon})R(t). \tag{6}$$

It is readily verified that, as required for a Laffer curve, $R(0) = C(0) = 0$ and $R(1) = C(1)$. If $\epsilon = 0$, $C(t) = 0$ as well. For $\epsilon > 0$, costs are positive for any $t > 0$. C is also an increasing function of the revenue share (that is, for our assumptions, it equals the tax rate t).

We assume net government revenue is redistributed in equal shares as a transfer—$T_i(t) = T(t) = N^{-1}\Sigma_i T_i(t)$. With this assumption and substituting the expressions for the governments' revenue (defined in Eq. (3)) and bureaucratic costs (defined in Eq. (6)) into the budget constraint (Eq. (2)), we find that the equal-size transfer $T(t)$ is given by

$$R(t) - \sum_i T_i(t) - C(t) = 0 \tag{7}$$

$$\sum_i T_i(t) = R(t) - C(T) \tag{8}$$

$$N^{-1} \sum_i T_i(t) = T(t) = N^{-1}(R(t) - C(T)) \tag{9}$$

[7] The assumption of full enforcement is not as restrictive as it might seem. In Appendix B, we show analytically for a general tax-avoidance function $f(t)$ that all model results are unaffected if agents' tax-avoidance decisions are homogeneous in the population. Intuitively, if all agents reduce their declared income by the same fraction f, both the transfer size and the tax payments for all agents are scaled by the same positive factor, leaving their voting decision unaffected. The empirical evidence indicates that tax compliance decreases with wealth levels (Alstadsæter et al. 2019) and that it is partially determined by social factors (Di Gioacchino and Fichera 2020). An interesting extension of the model could thus be to integrate income-dependent tax-compliance decisions that are mediated by average compliance in the respective ego network.

$$= N^{-1}(Nt\bar{y} - (1-(1-t)^{\epsilon})Nt\bar{y}) \tag{10}$$

$$= t(1-t)^{\epsilon}\bar{y}. \tag{11}$$

The above equation directly implies the existence of a Laffer curve for net tax revenues (gross revenues minus costs) and, consequently, transfer sizes, as also shown in Fig. 3. It is straightforward to calculate this revenue-maximizing tax rate as $t^* = 1/(1+\epsilon)$. This revenue-maximizing rate approximates the optimal rate of the poorest agent(s), whose wages are close to zero, as (almost all) of their income comes from the transfer.[8] In this sense, the transfer-maximizing rate approximates the rate chosen using a Rawlsian maximin criterion (Rawls 1974).

As the resulting revenue is then equally split among all agents, this implies the following difference between pre-tax and post-tax income for each agent *i*:

$$\Delta y_i = -t \cdot y_i + t \cdot (1-t)^{\epsilon} \cdot \bar{y} \tag{12}$$

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 income for $\epsilon = 0$ and shrinks in ϵ and *t* for $\epsilon > 0$.

As a corollary of the assumption of lump-sum transfers at the macro level, taxation always decreases wage inequality. To make this precise, consider the ratio between the pre- and post-tax Gini coefficients—*G* and $\tilde{G}$, respectively. This ratio can be expressed analytically as

$$\frac{\tilde{G}}{G} = \frac{1-t}{(1-t) + t \cdot (1-t)^{\epsilon}}. \tag{13}$$

It follows that $\tilde{G}/G < 1$ if $t > 0$ and ϵ is finite. The inequality-decreasing effect is mitigated for a larger ϵ for any given *t* since higher taxation inefficiency decreases the size of the

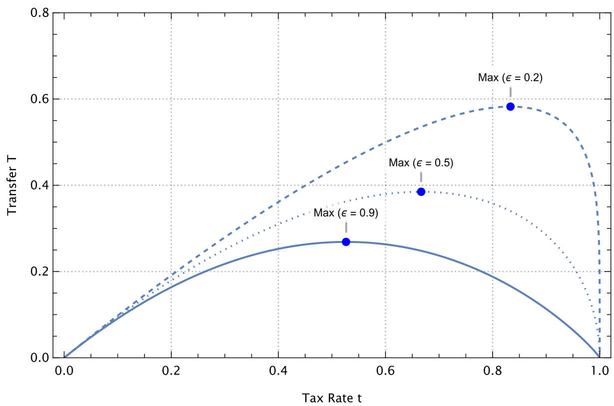

Fig. 3 Laffer curves for different levels of $\epsilon \in \{0.2; 0.5; 0.9\}$, representing taxation inefficiency resulting from increasingly necessary bureaucratic efforts. Laffer curves feature a unique revenue-maximizing rate that is further to the right the lower ϵ is

[8] For an exponential pre-tax 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*.

(equalizing) transfer. As a limit case, inequality is unchanged for any tax rate if ϵ approaches infinity—that is, $\lim_{\epsilon\to\infty} \tilde{G}/G = 1$—with the transfer then being nil.

3.5 Localized individual perceptions and tax-rate acceptance

In considering decisions about whether to accept any given tax rate *t*, we assume agents are purely selfish—that is, 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 elasticity ϵ in the static setting. However, agents' sensing is imperfect because of the echo chambers that emerge from the interplay of income inequality and homophilic segregation (see the middle and right panels of Fig. 2). Thus, biased perceptions result from limited epistemic diversity, as discussed in detail in Guerrero (2021). In particular, we assume their perception of the mean gross income to be a compound of the actual global mean $\bar{y}$ and the mean income in their ego network l_i with a parameter $a \in [0, 1]$.[9] Therefore, the lower *a*, the stronger the effect of echo chambers. Hence, agent *i* believes the mean income to be $\hat{y}_i$—that is,

$$\hat{y}_i = a \cdot \bar{y} + (1 - a)l_i. \tag{14}$$

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

$$y_i^* = (1 - t)^\epsilon \cdot [a \cdot \bar{y} + (1 - a) \cdot l_i] \tag{15}$$

The weight parameter $a \in [0, 1]$ is identical for all agents; $a = 1$ indicates 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; if $l_i > \bar{y}$, they accept some tax rates, meaning a net loss.

3.6 Voting behavior 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 on any given tax rate *t*. Namely, agent *i* with pre-tax income y_i follows the voting rule V_i:

$$V_i = \begin{cases} 1\text{if } (1 - t)^\epsilon \cdot [a \cdot \bar{y} + (1 - a) \cdot l_i] > y_i \\ 0\text{if } (1 - t)^\epsilon \cdot [a \cdot \bar{y} + (1 - a) \cdot l_i] = y_i \\ -1\text{if } (1 - t)^\epsilon \cdot [a \cdot \bar{y} + (1 - a) \cdot l_i] < y_i \end{cases} \tag{16}$$

[9] 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 downward by a factor of $(1 - a)$—that is, individual perceptions approach the true mean income $\bar{y}$ linearly in *a*.

The individual voting decision of any agent regarding a tax rate $0 < t < 1$ is therefore a function of their perception l_i, the elasticity parameter ϵ, their pre-tax 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 (that is, without misperceptions). It follows that an individual tax rate *t* has a majority if $V > 0$[10]:

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

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

3.7 Dynamic updating and beliefs about efficiency and local average income

So far, we have only considered static redistribution and voting behavior in one period. However, it is perhaps not reasonable to assume that misperceptions fully persist, even though realized transfers should show agents that their beliefs must be wrong. To account for this, we introduce dynamic updating via adaptive expectations for agents in two learning scenarios: for the efficiency parameter ϵ_i and for perceived local income l_i. The main idea with these adaptive expectations is that agents form initial beliefs about the efficiency parameter or the local income and revise them as new evidence appears. If the new observation is higher than the previous-period expectation, they revise their estimate upward, and if it is lower, they revise it downward.

The dynamic extension for ϵ is motivated by the empirical finding that efficiency preferences mediate the relationship between income position and redistribution (Tepe et al. 2021). The extension then allows us to examine which types of initial states lead to *learnable* unbiased tax rates—that is, collective decisions that converge to the MR benchmark with $a = 1$. We assume agents correctly observe the transfer *T* and their own income y_i but sustain their misperception about societal average 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 $\epsilon^e_{i,\tau}$ for time $\tau = 0, 1, \ldots$, so

$$\epsilon^e_{i,\tau} = \epsilon^e_{i,\tau-1} + \lambda(\epsilon_{i,\tau-1} - \epsilon^e_{i,\tau-1}). \tag{17}$$

[10] The aggregation function assumes full voter turnout. A potentially interesting extension would be to incorporate vote abstention—for example, abstention resulting from the closeness of elections, which theoretically biases collective outcomes (Landgren et al. 2021). Empirically, the effect appears ambiguous, though (Sabet 2023).

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 in the previous period (Evans and Honkapohja 2001). Rewriting this equation and substituting the expression for the perceived realization of ϵ_i into Eq. (17) yields the updating rule

$$\epsilon^e_{i,\tau} = \lambda \cdot \frac{\log\left(\frac{t_{\tau-1}y_i+T_{\tau-1}}{t_{\tau-1}(a\cdot\bar{y}+(1-a)l_i)}\right)}{\log(1-t_{\tau-1})} + (1-\lambda)\cdot\epsilon^e_{i,\tau-1}. \tag{18}$$

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

In the second dynamic extension, we consider the polar case and let agents adjust their estimate of local mean income l_i while keeping their beliefs about ϵ constant (and correct). Here, we follow the intuition that agents might not have strongly biased opinions about the deadweight loss of taxation as reflected in ϵ but might correct their belief about the level of taxable income that is to be redistributed.[11] This estimate can also be deduced from the expression for the transfer T and yields, similarly, $l_i = (a\bar{y} - \frac{(1-t)^{-\epsilon}(ty+T)}{t})/(a-1)$. Using the same type of updating rule, we can express the learning mechanism for the perceived local income l^e_i at time τ as

$$l^e_{i,\tau} = l^e_{i,\tau-1} + \lambda(l_{i,\tau-1} - l^e_{i,\tau-1}). \tag{19}$$

Substituting again for the realized value of local mean income l, we can express the corresponding updating mechanism for agent i as

$$l^e_{i,\tau} = \lambda \cdot \frac{a\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i+T_{\tau-1})}{t_{\tau-1}}}{a-1} + (1-\lambda)l^e_{i,\tau-1}. \tag{20}$$

We now examine the effects of both updating rules on voting behavior.

4 Results

4.1 Simulation mode and validation

To ensure internal validity, we simulate 100 Monte Carlo runs for each parameter combination and report simulation averages, if not otherwise stated. The homophilic linkage mechanism as the central force in our model is also externally valid following Windrum et al. (2007), in terms of both input calibration (income distribution, number of close contacts, and homophilic linkage, as discussed above) and reproduction of key empirical observations in its output. One such observation is that the model mechanism yields perception patterns of income inequality in general (Schulz et al. 2022b) and of wage gaps (Mayerhoffer and Schulz 2022).

[11] We thank Michael Neugart for pointing us to this extension and the underlying intuition.

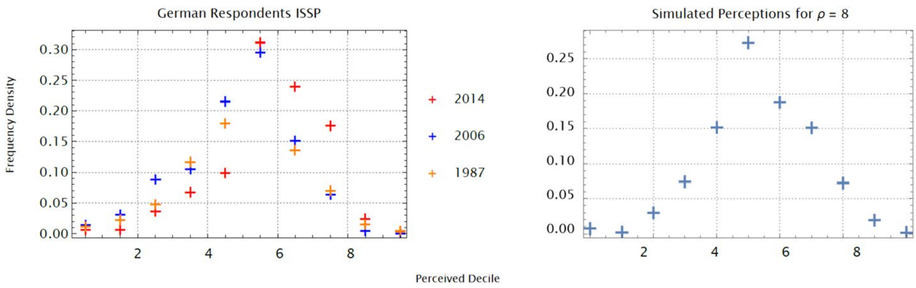

Fig. 4 The figure shows self-perceptions of income deciles from a survey for German respondents in the International Social Survey Program (left panel) and as simulation outcomes in the right panel for $\rho = 8$. The middle-class bias that emerges in the simulation closely mimics the middle-class bias that empirical surveys consistently demonstrate

For sufficiently high levels of inequality and homophilic linkage, agents tend to have the median income in their ego network, in line with empirical studies finding a "middle-class bias" (Knell and Stix 2020; Schulz et al. 2022b). To illustrate this point, in Fig. 4 we show the results for a simulated network with homophily parameter $\rho = 8$ and an exponential pre-tax income distribution alongside empirical self-perceptions in Germany gathered from the International Social Survey Program. The network simulation replicates the empirical findings remarkably well and should thus constitute a good starting point to model real-world perceptions.[12] Intuitively, this feature of the model is a direct corollary of the homophilic segregation in the model: Since the weighting function is inversely related to the distance in income, for any $\rho > 0$ it will assign the highest weight to the set of neighbors for an agent i to whom the income distances are minimal. For an exponential distribution, these are always the incomes of the agents that are distributed symmetrically around agent i, except for corner cases close to the minimum and maximum observation. Schulz et al. (2022b) prove this intuitive result analytically. This also directly implies that an increase in income inequality, measured by an increasing mean-to-median ratio, is not reflected in the agents' perceptions within their ego networks.

We use the latter fact below. The first results concern the static setting of the model—that is, agents not updating their beliefs—and we go on to present the results for belief adaptation in a dynamic setting later in this section.

4.2 Acceptance of tax rates depending on the weight of global signal

Empirically, perceived levels of inequality, not actual levels, drive redistributive behavior (Choi 2019). We examine this issue by simulating various a's that represent the quality of agents' information regarding transfer size. Figure 5 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 localized perception for the assumption of an exponential pre-tax income distribution. For $a = 1$, our model nests the MR model with perfect sensing. The plot line for $a = 1$ can be understood as 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 redistribu-

[12] Network-generated misperceptions do not require behavior that is implausibly persistently irrational without any adaptation. By contrast, agents in our model rationally extrapolate from skewed information sets.

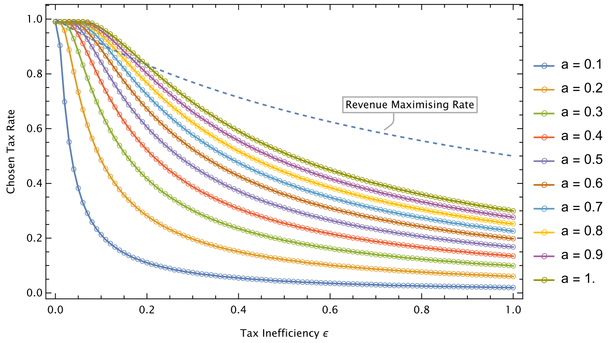

Fig. 5 The figure shows the implemented tax rates both for different weights for the global signal $a \in [0,1]$and for varying the elasticity of taxable income $\epsilon \in [0,1]$. The simulation holds the homophily level constant at $\rho = 8$. Increasing the weight of the global signal and improving the accuracy of perceptions unanimously increases (implemented) redistribution since agents then expect higher transfers on average

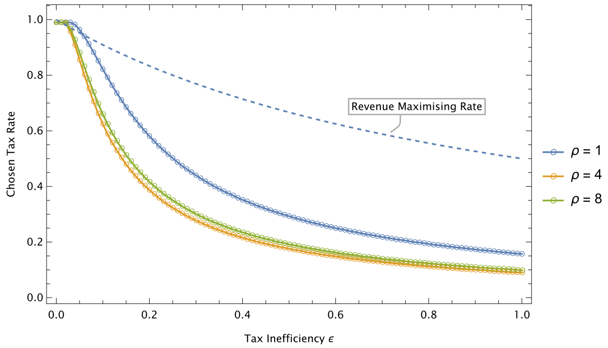

Fig. 6 The figure shows the implemented tax rates for varying the elasticity of taxable income $\epsilon \in [0,1]$. For the simulation, the weight of the global signals is held constant at $a = 0.3$, and the homophily strength varies in the range $\rho \in \{1; 4; 8\}$. Increasing network segregation decreases (implemented) redistribution

tion since the mean income in their ego networks is lower than the global mean. Thus, the accepted tax rate decreases as the weight put on the localized perception increases (that is, when a is lower).[13] We can thus replicate the finding in the literature (for example, in Choi (2019)) that misperceptions constitute a good, albeit partial, candidate explanation for low redistribution in reality.

4.3 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 Fig. 6 highlights, again for an exponential pre-tax income distribution. 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 greater 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 perceived mean remains nearly unaffected. Our model thus replicates the finding that segregation tends to

[13] That for low inefficiency levels, the highest tax rate capable of winning a majority exceeds the revenue-maximizing 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.

lower (implemented) redistribution (Franko and Livingston 2022) and explicates the channel—namely, segregation leads the poorer part of the population to underestimate the size of a potential transfer and the benefits from redistribution.

4.4 Redistribution and inequality with biased beliefs

So far, we have only considered redistribution in a setting in which the level of inequality is fixed since we have only considered an exponential pre-tax income distribution with a fixed Gini of $G \approx 0.5$. A central empirical finding, however, is that redistribution does not appear to react strongly to actual inequality and that the relationship is much stronger for perceived inequality. To examine how beliefs biased by homophilic segregation mediate the relationship between (pre-tax) inequality and redistribution, we use a lognormal distribution and vary its dispersion parameter $\sigma \in [0.01, 1]$ in increments of 1% to see how inequality affects realized redistribution—that is, the chosen tax rate t.[14] Figure 7 shows the relationship between the pre-tax Gini coefficient and redistribution for various $a \in [0, 1]$. A weight of $a = 1$ again corresponds to the MR benchmark, and with decreasing a, agents put more emphasis on their local, biased perceptions.

Figure 7 shows that biased perceptions due to homophilic segregation can indeed mediate and weaken the link between actual inequality and redistribution. The full-information case corresponding to the MR model is given for $a = 1$, in which agents only rely on the unbiased global signal to estimate the societal mean income. The lower a is, the higher the effect of the agents' biased perception sets. The bias results from homophilic segregation: While the *actual* mean-to-median ratio rises unambiguously with σ and the Gini coefficient of the pre-tax distribution, the ratio of *perceived* mean incomes to the agents' own incomes increases much slower because of network segregation and because their perception sets are unrepresentative of the whole distribution. This is ultimately the result of the interaction of (pre-tax) inequality and homophilic linkage that we discussed in Sect. 3. Since segregation

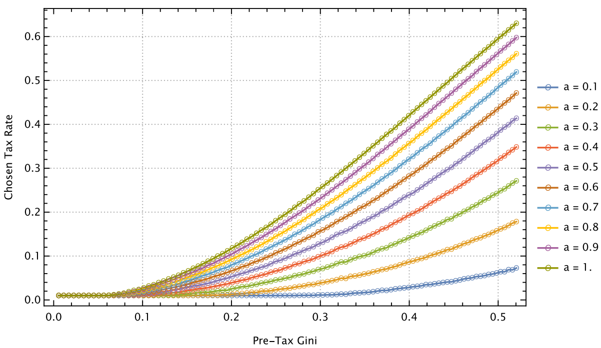

Fig. 7 The figure shows the (average) implemented tax rates for different weights for the global signal $a \in [0; 1]$ and for varying the Gini coefficient of the pre-tax income distribution that is initialized as a lognormal distribution. The simulation holds the homophily level constant at $\rho = 8$. A stronger reliance of agents on their local signal (a lower a) has a mediating effect on the relationship of pre-tax inequality and the chosen tax rate t, with the relationship weakening the lower a is

[14] The Gini coefficient for any lognormal distribution is given by $G = \operatorname{erf}(\sigma/2)$.

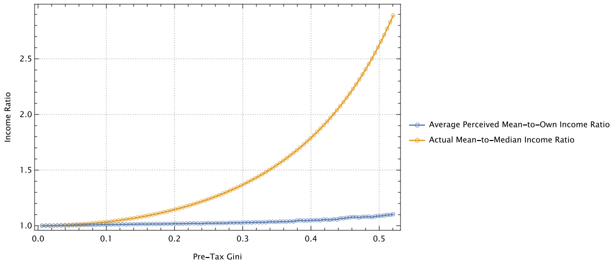

Fig. 8 The figure shows the impact of an increase in (pre-tax) inequality for a lognormal income distribution on both the mean-to-median income ratio and the perceived mean-to-own income ratio for perceptions based on a homophilic network with $\rho = 8$ and averaged over all agents and MC runs. As expected, the mean-to-median ratio increases rapidly with increasing inequality, but the perceived ratio increases much slower in the pre-tax Gini. This is because increasing inequality also increases network segregation and perception sets, thus growing more skewed with an increasing pre-tax Gini coefficient

also increases with (pre-tax) inequality, the effect of an increasing mean-to-median ratio is dampened for the perception sets.

Figure 8 shows the behavior of the actual and perceived mean-to-median ratios for $\rho = 8$. As is immediately visible, there exists a gap between correct and biased perceptions that is growing in pre-tax inequality. This gap is also reflected in the chosen tax rates shown in Fig. 7: For low levels of pre-tax inequality, biased and unbiased perceptions align, and so do the chosen tax rates. As biased and unbiased perceptions diverge, the implemented tax rates for biased perceptions also diverge from the full-information case with $a = 1$. The increase in σ and the increase in pre-tax inequality is not reflected in an increased average preference for redistribution. Segregation-induced biased perceptions thus constitute a candidate explanation for the empirically dampened response of redistributive preferences to increases in pre-tax inequality, especially if pre-tax inequality is already high.

4.5 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, as we show for our two dynamic extensions of the static model.

In the first extension, we let agents adapt their beliefs about taxation efficiency ϵ over time according to the updating rule in Eq. (18). Beliefs about ϵ in the first period are always initialized as correct—that is, $\epsilon_{i,0} = \epsilon \ \forall i$. This assumption appears sensible to us, as it amounts to saying agents know how vigorously they will attempt to hide their income for a given tax rate ϵ, initially (correctly) extrapolating to the whole distribution and only updating their beliefs when they are confronted with opposing evidence.[15] Figure 9 shows three exemplary belief schedules ϵ on the y-axis against pre-tax incomes y_i on the x-axis 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 being convinced that redistribution increases the aggregate income to be redistributed

[15] The qualitative model dynamics reported below are seemingly robust for other initial values, though. Results for these robustness checks are available upon request.

Fig. 9 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 evident in the belief schedules' growing slope over time

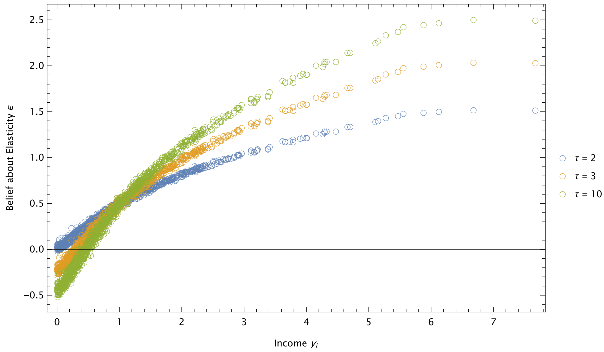

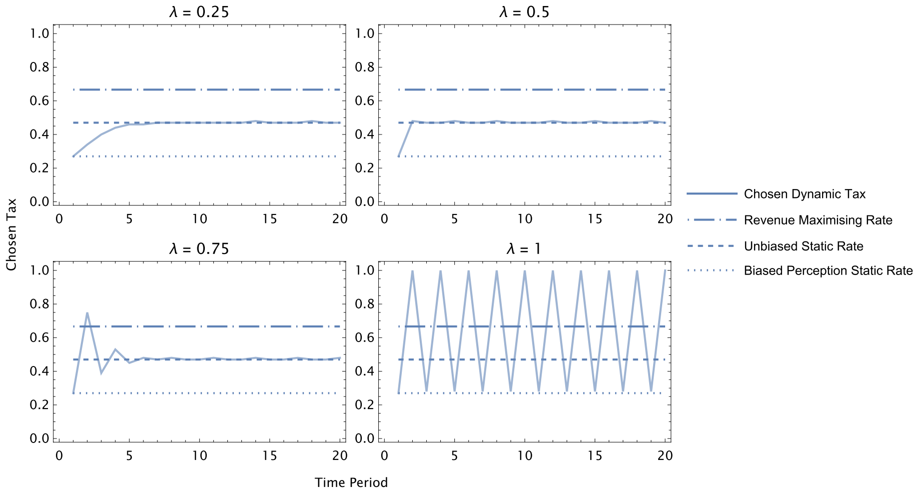

Fig. 10 The figure shows the chosen tax rates for the dynamic updating process for ϵ with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$. All simulations are conducted for $a = 0.5$, $\rho = 8$, and a true $\epsilon = 0.5$. The upper two panels show monotonic convergence, and the lower left panel shows oscillatory convergence to the unbiased static rate of $a = 1$. By contrast, the lower right panel shows a two-cycle of tax rates that constantly alternate between unity and the static biased-perception rate. The displayed patterns remain stable for longer periods

($\epsilon < 0$). 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 Fig. 10). The mechanism for this opinion polarization is 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 realized 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 polarized: 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, because of the effects of localized perception, the poorer individuals were likely to not 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. Consequently, the polarization in perceptions of tax efficiency leads to a growing accord about acceptable tax rates; hence, the polarizing views on efficiency actually drive the convergence to a steady state for the implemented tax level.

In our model, in the steady state (which is already reached for $\tau = 10$ in Fig. 9), almost all agents hold wrong beliefs about both societal mean income and 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 behavioral literature on the topic (see, for example, Caplan, 2001 on so-called "rational irrationality").

Note that poor agents' apparently absurd belief that $\epsilon_i < 0$ is inconsequential for our simulation results. In this situation, agents believe that the sum of taxed money somehow increases whenever it is given out by the state as a transfer, which indeed makes little economic sense. Yet we can impose $\epsilon_{i,\tau} \geq 0 \ \forall \ i, \tau$ without any impact on the chosen tax rates, as we also verify in our simulations.[16] The reason for this is that all agents with pre-tax income below the true mean income will rationally accept any tax rate $t > 0$ whenever they believe ϵ to be zero (Foley 1967). The voting decision of these agents will thus be exactly the same, irrespective of whether ϵ_i is bounded at zero.

The discussed perceived elasticities give rise to three qualitative types of emergent voting behavior: monotonic convergence, oscillatory convergence, and limit cycles. These three ideal cases are illustrated for a particular parameter constellation in Fig. 10—that is, for $a = 0.5$, $\rho = 8$, and a true $\epsilon = 0.5$, with several values for the error-correction parameter $\lambda \in \{0.25; 0.5; 0.75; 1\}$. Our results can thus be interpreted as capturing ideal types of belief updating behind the empirical variations in the demand for redistribution (Georgiadis and Manning 2012).[17]

Monotonic convergence is illustrated in Fig. 10 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 downward, leading them to vote for higher tax rates t. Since beliefs are only partially adjusted because of a comparatively low λ, this self-reinforcing process of higher-than-predicted transfers, a downward adjustment of beliefs about ϵ, and acceptance of 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.[18]

The same basic mechanism is behind oscillatory convergence, where the 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 expected and adjust their beliefs accordingly. Nevertheless, since λ is higher, the adjustment is stronger and leads them to collectively vote for a tax rate above the unbiased static rate. For this new tax rate, the generated actual inefficiency is rather high because of the high chosen tax

[16] A summary of those results for a given parameter constellation can be found in Appendix D.

[17] We always only show 20 periods to enhance visibility. Identical patterns appear to continue for a much higher number of periods—for example, 100.

[18] The tax rate still exhibits tiny fluctuations of around 1 percentage point. This artifact results from the voting procedure that only allows agents vote for tax rates in increments of 1%.

rate, leading poorer agents to experience lower-than-expected transfer sizes and to adjust their beliefs about the tax-efficiency parameter ϵ upward 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-maximizing rate, approximating the rate chosen under a Rawlsian maximin criterion, does not emerge in any of our static or dynamic setups. Moreover, even if it did, it would do so only accidentally for very particular parameter constellations and not because of a general convergence mechanism. This might imply that while the unbiased static rate is learnable even under misperception, the revenue-maximizing rate is not. To let agents self-organize around the Rawlsian maximin rule thus necessitates a different voting regime.

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—that is, a tax rate of (almost) unity.[19] 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 (Aguiar-Conraria et al. 2012). Electoral cycles are thus not necessarily the result of irrational bandwagon effects or an irrational taste for change in the electorate; they can also emerge from rational learning. However, the collective voting behavior only converges, even for low values of λ (either oscillatory or monotonic), whenever initial beliefs do not deviate too far from the unbiased benchmark, as we show in Appendix C. Importantly, homophilic segregation thus not only causes misperceptions in a static setting and hinders redistribution but also can prevent agents from learning dynamically and self-organizing around the unbiased rate.

In the second extension, we let agents update their (initially biased and individualized) beliefs about local mean income. Using the compound rule in Eq. (14), the belief of an agent i is thus given as $\tilde{y}_{i,\tau} = a\bar{y} + (1-a)l^e_{i,\tau}$. As with the first scenario, we initialize the simulations with $l^e_{i,0} = l_i \; \forall i$ for the updating rule to be well defined.

The resulting dynamics are much less varied for this type of updating compared to the updating process for ϵ. In all considered cases in Fig. 11, the chosen tax rates converge monotonically to the unbiased static rate. This finding is hardly surprising. While the updating of ϵ leads to a situation in which wrong beliefs about ϵ might (or might not) compensate

[19] For computational reasons, we set an upper limit of $t = 0.9999$. The model shows the same dynamics for different upper bounds close to one.

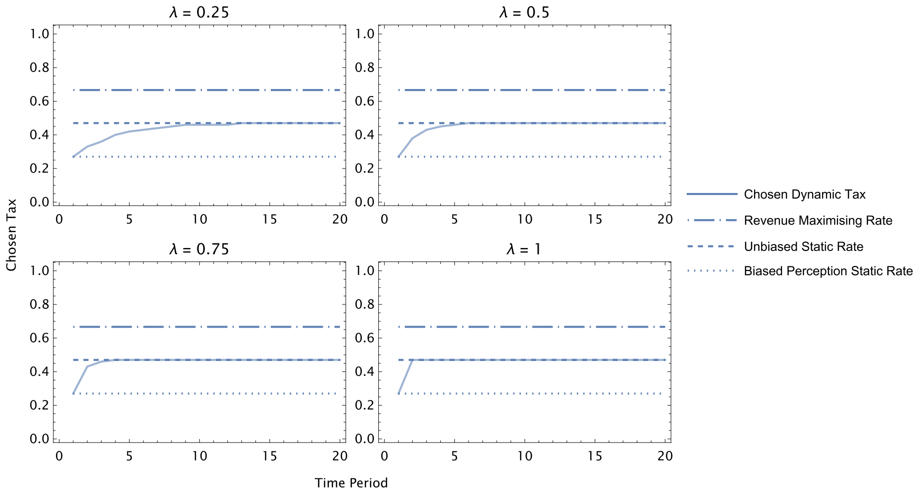

Fig. 11 The figure shows the chosen tax rates for the dynamic updating process with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$ for the updating of l. All simulations are conducted for $a = 0.5$, $\rho = 8$, and $\epsilon = 0.5$. In all cases, the chosen tax rates converge monotonically to the unbiased static rate, as beliefs about l get closer to the true value over time

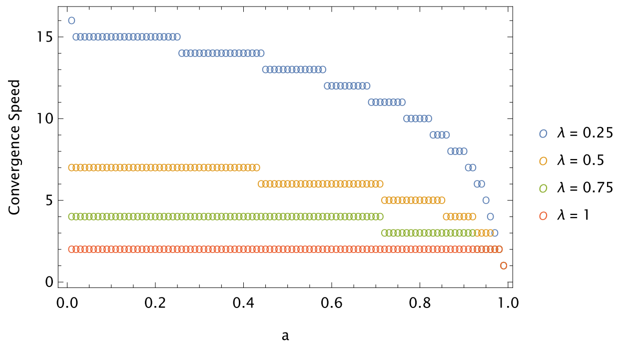

Fig. 12 The figure shows the speed of convergence for different values of the weight for the global signal $a \in [0, 1]$ for the chosen rate to reach the unbiased static rate in voting behavior for the updating rule for the local average income l. All simulations are carried out with $\rho = 8$, $\epsilon = 0.5$, and $\lambda \in \{0.25; 0.5; 0.75; 1\}$. The convergence speed increases when a is higher—that is, the more weight agents put on the global signal and the less biased their initial beliefs are and in the intensity of error adjustment λ .

for wrong beliefs about the level of societal income, learning about l leads to less biased estimates of l and, consequently, the societal average income $\bar{y}$. In the second scenario, values of l_i thus converge to $\bar{y}$, which implies that perceptions of average societal income $\hat{y}_i \to \bar{y} \; \forall \; i$ as the unbiased value. In our model context, oscillations and oscillatory convergence are thus only consistent with updating ϵ, not updating l.

Yet the speed of convergence to the unbiased static rate obviously depends on initial beliefs and the intensity of error adjustment captured by λ. Figure 12 shows the speeds of convergence, defined as the number of periods until the chosen tax rate has reached the unbiased static rate, for $a \in [0, 1]$ and various $\lambda \in \{0.25; 0.5; 0.75; 1\}$. The higher the a, the lower the number of periods until the unbiased state is reached. Thus, the more biased initial beliefs are, the more adaptive learning about l is necessary to mitigate this bias. Similarly, the stronger the error adjustment, the faster the convergence, as is obvious from the convergence speeds for differing λ's in Fig. 12.

A comparison of the two learning regimes shows that they yield two contrasting conclusions about the impact of the error-adjustment parameter λ. While a higher λ tends to destabilize the system and leads to oscillatory behavior for adaptive learning of ϵ, it unambiguously leads to faster convergence to the unbiased static rate, if agents update their beliefs about the local average income l_i. The empirical literature indeed suggests that electoral cycles exist (see Aguiar-Conraria et al. 2012, for a review). As we showed, only the belief updating about taxation efficiency ϵ is capable of generating these cycles in our modeling context. If agents updated on their beliefs about societal income, our model would predict no cyclical behavior. Thus, empirical tests about which types of learning are empirically prevalent might help to determine whether our stylized explanation of electoral cycles is valid.

5 Discussion

This study presented a model that suggests that income misperceptions and, consequently, paradoxical voting behavior can result from correct belief formation, given segregation in social networks. The model is based on a parsimonious representation of taxation and voting behavior, and it uses a Random Geometric Graph–type model to represent the observed homophily. The results of the model support the hypothesis that individuals' 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 society as a whole is.

The model results support our hypothesis that misperceptions of average societal income that are caused by segregation in social networks tend to decrease redistribution, as shown in a simulation of an exponential initial (labor) income distribution that approximates empirical distributions well. The high level of segregation in the network structure, coupled with the agents' bias toward assuming their own income rank and income to be close to the middle, leads to a significant underestimation of societal income. This has important implications for the relationship between taxation and revenue, as optimal tax rates and revenue depend on agents' perceptions of overall mean income. Regarding our opening question on the link between redistribution and *increasing* pre-tax inequality, we also find that misperceptions due to homophilic segregation tend to mediate and weaken the link. The more agents rely on the biased signal from their local social networks, the weaker the relationship between pre-tax inequality and redistribution, qualitatively replicating the findings in, for example, Choi (2019).

Extending our model to a dynamic setting by allowing for the endogenous adjustment of beliefs regarding taxation efficiency led to a rich set of voting dynamics. We showed that the beliefs of poor people regarding taxation efficiency tend to be optimistic, while the rich are more pessimistic. Suppose error correction is slow or network segregation and perceptions of mean societal income are close to the true mean. Then voting behavior converges to the unbiased tax rate for correct initial perceptions. Misperceptions can thus be overcome under certain conditions. For rapid error correction, cyclical voting behavior emerges because of endogenous waves of optimism and pessimism regarding taxation efficiency that lead agents to over- and undershoot the unbiased rate. Thus, our model might also help to explain empirical electoral cycles by reference to endogenous changes in the overall degree of optimism. By contrast, if agents update their beliefs about societal income, error adjustment unanimously leads to monotonic convergence to the unbiased state. Thus, our model suggests that electoral cycles are only consistent with learning about efficiency, not learning about societal incomes. In the latter case, rapid error adjustment increases the speed of convergence to the unbiased state, in contrast to the case of the learning-about-efficiency mechanism, in which the intensity of error adjustment has a destabilizing effect. With these opposing conclusions, we hope to invite research on which of the two learning regimes dominates empirically or whether the empirically documented empirical cycles are caused by other factors.

Even though our results rely on rather specific assumptions regarding the Laffer curve, income distribution, perception and expectation formation, voter turnout, and aggregation mechanisms, they provide a robust joint explanation of why perceptions are so skewed and how they might mediate the nexus between inequality and redistribution. Yet, given our simulation approach, this explanation is not a "how-actually" explanation but a "how-possibly," or candidate, explanation (Grüne-Yanoff and Verreault-Julien 2021), and there might be alternative candidates. The model serves as a surrogate system (Mäki 2009), and given its internal and external validity as explained in Sect. 4, the surrogate resembles its target system: real-world societies. Consequently, the how-possibly explanation given in this paper goes beyond the criteria raised by Grüne-Yanoff and Verreault-Julien (2021) for a successful epistemically possible how-possibly explanation: Those criteria only demand that the explanatory mechanism not be ruled out by epistemic reality, whereas in the present case, known empirical facts and the proposed mechanism mutually support each other. As a next step toward a how-actually explanation, researchers need to look for empirical (counter)evidence for the proposed mechanism.

Much like in our model, (perceived) self-interest is a crucial determinant of real-world voting behavior in redistributive settings, yet this effect is mediated by efficiency concerns that counteract immediate perceived 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 it might create (Berens and von Schiller 2017; Atria 2023). 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 view of the "leaky bucket" view (Okun 1975) of redistribution, believing that enforcing taxation is costly and leads to deadweight loss. By contrast, poorer people appear to converge to the belief that the state is very efficient at redistribution and that taxation comes with hardly any costs, if at all. Rich people's opposition to redistribution, which is well documented in the survey by Arndt (2020), might, therefore, be caused not (only) by self-interest, but (also) by biased beliefs that themselves are the result of correct belief updating from a biased signal.

Regarding this channel, to our knowledge, only three works have investigated how the perception of the degree of government waste varies with income: Two early contributions by Hasenfeld and Rafferty (1989) and Brooks and Brady (1999) show for the US that all measures of government efficiency are negatively correlated with income, in line with our hypothesized mechanism. By contrast, Williamson (2019) shows a positive correlation between income and perceptions of governmental efficiency. Yet she also demonstrates that citizens' concept of "government waste" is quite far from the expert notion and includes all disliked policies rather than only "true" inefficiencies. Given the potential relevance of perceived efficiency for our dynamic model, there thus exists a possible avenue for further research to examine how the opinion polarization of our model relates to the actual affective polarization by socioeconomic status.

With our theoretical model, we have explicated several possible mechanisms regarding the nexus of network segregation, perceptions of inequality and societal average income, and redistributive taxation. While extensive empirical exercises are beyond the scope of this paper, we hope to invite empirical research on several questions our model sheds light on: Does network segregation indeed respond to income inequality, and does it decrease redistributive preferences and taxation? On which dimension do people tend to update their beliefs—average income, taxation efficiency, or something else entirely? Since cyclical behavior can only emerge in our model if agents update their views of ϵ, determining whether this type of updating happens in reality can also help to determine whether our attempt at explaining electoral cycles is sensible and to discriminate between contrasting explanations. Are the involved timescales in reality short enough to enable convergence to the unbiased state? And do the rich indeed have less favorable views about the efficiency of redistribution than the poor?

Model extensions could include the role of the media, education, or political polarization to account for parameter heterogeneity that we did not explicitly model. Furthermore, the model could be adapted to other voting and tax regimes, including progressive ones, to examine the robustness of the results. The empirical literature has, for example, identified the electoral calendar and the political regime as relevant determinants of political business cycles (Dubois 2016). While our perceptions based on homophilic networks appear externally valid, belief formation lacks a purposive element—for example, regarding communication in networks. With our model, we hope to have provided a first benchmark that can be extended along those lines.

5.1 Supplementary information

The network model to reproduce the data used in this paper can be found in Mayerhoffer and Schulz (2025).

Appendix A

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 $\bar{y}$ as the vector of incomes and $\bar{p}$ as the vector of perceived incomes. Finally, normalize 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

$$B\vec{y} = \vec{p}. \tag{A1}$$

Writing this relationship as a sum, we get

$$l_i = \sum_{j=1}^{N} b_{ij} \cdot y_j. \tag{A2}$$

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

$$l_i = \bar{y} + (N-1) \cdot cov(b_{ij}, y_j). \tag{A3}$$

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

$$\hat{y}_i = a \cdot \bar{y} + (1-a)l_i. \tag{A4}$$

Substituting l_i from eq. (A3) yields

$$\hat{y}_i = a \cdot \bar{y} + (1-a)(\bar{y} + (N-1) \cdot cov(b_{ij}, y_j)) \tag{A5}$$

and rearranging finally implies

$$\hat{y}_i = \bar{y} + (1-a) \cdot (N-1) \cdot cov(b_{ij}, y_j). \tag{A6}$$

Compared to the purely localized perceptions given in Eq. (A3), the compound rule in Eq. (A3) thus scales the bias term down by $(1-a)$.

Appendix B

Successful homogeneous tax avoidance

The basic model setup presented in Sect. 3 assumes that all tax avoidance efforts of agents are unsuccessful. Yet, it appears implausible that agents would engage in avoidance efforts if they continuously are not successful. In this appendix, we show that homogeneous avoidance responses do not affect the results presented in the main body of the paper at all. We assume that all agents declare only a fraction $1-f(t)$ as their taxable income, with f as the fraction of avoided taxation. By the discussion above, $f'(t)>0$ and $0\leq f(t)<1,\ \forall t$. To simplify notation, denote any variable $x(f)$ as affected by tax avoidance with f as the fraction of undeclared income. Perceptions are formed for the declared incomes, which implies that $l_i(f)=L_i^{-1}\Sigma_{j\in\Omega_i}(1-f(t))y_j=(1-f(t))l_i$, with L as the number of agents in agent i's ego-network and Ω_i as the set of those agents. $y_i(f)$, $\bar{y}(f)$ and $l_i(f)$ are thus just rescaled versions of y_i, $\bar{y}$ and l_i, respectively. Substituting these into the voting decision of agent i in Eq. (16) shows that this rescaling does not alter their decision in any way.

Starting from the decision rule with adjusted incomes as

$$V_i(f)=\begin{cases} 1\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]>y_i(f)\\ 0\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]=y_i(f)\\ -1\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]<y_i(f),\end{cases} \tag{A7}$$

we can substitute the expressions for $y_i(f), \bar{y}(f)$ and $l_i(f)$ into the decision rule and find that the rule is equivalent to the one based on full tax enforcement V_i in Eq. (16) by noting that $(1-t)^{\epsilon}\cdot[a\cdot(1-f(t))\bar{y}+(1-a)\cdot(1-f(t))l_i]=(1-t)^{\epsilon}\cdot(1-f(t))\cdot[a\cdot\bar{y}+(1-a)\cdot l_i]$ as

$$V_i(f)=\begin{cases} 1\text{if } (1-t)^{\epsilon}\cdot(1-f(t))\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]>(1-f(t))y_i\\ 0\text{if } (1-t)^{\epsilon}\cdot(1-f(t))\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]=(1-f(t))y_i\\ -1\text{if } (1-t)^{\epsilon}\cdot(1-f(t))\cdot[a\cdot\bar{y}(f)+(1-a)\cdot l_i(f)]<(1-f(t))y_i,\end{cases} \tag{A8}$$

and letting $(1-f(t))$ cancel out on both sides of the (in)equalities as

$$V_i(f)=V_i=\begin{cases} 1\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}+(1-a)\cdot l_i]>y_i\\ 0\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}+(1-a)\cdot l_i]=y_i\\ -1\text{if } (1-t)^{\epsilon}\cdot[a\cdot\bar{y}+(1-a)\cdot l_i]<y_i.\end{cases} \tag{A9}$$

Thus, for homogeneous avoidance of all agents (possibly as a function of the tax rate t), voting decisions are equivalent to the case with unsuccessful avoidance.

For belief-updating, we also find that voting decisions under updating of ϵ and l_i are invariant to rescaling incomes by $(1-f(t))$. Consider first the case of ϵ. The belief about ϵ at τ, $\epsilon^e_{i,\tau}(f)$, is given by

$$\epsilon^e_{i,\tau}(f)=\lambda\cdot\frac{\log\left(\frac{t_{\tau-1}y_i(f)+T_{\tau-1}(f)}{t_{\tau-1}(a\cdot\bar{y}(f)+(1-a)l_i(f))}\right)}{\log(1-t_{\tau-1})})+(1-\lambda)\cdot\epsilon^e_{i,\tau-1}. \tag{A10}$$

From the expression for the actual transfer, we find that $T_{\tau-1}(f) = t_{\tau-1} \cdot (1-t)^{\epsilon} \cdot \bar{y}_i(f) - t \cdot y_i(f) = t_{\tau-1} \cdot (1-t)^{\epsilon} \cdot (1-f(t))\bar{y}_i - t \cdot (1-f(t))y_i = (1-f(t))T_{\tau-1}$. Unsurprisingly, the transfer size is therefore also scaled by the same factor. Substituting the expressions for $T_{\tau-1}(f), \bar{y}(f), y_i(f)$ and $l_i(f)$ into Eq. (B10) yields

$$\epsilon^e_{i,\tau}(f) = \lambda \cdot \frac{\log\left(\frac{t_{\tau-1}(1-f(t))y_i + (1-f(t))T_{\tau-1}}{t_{\tau-1}(a\cdot(1-f(t))\bar{y} + (1-a)(1-f(t))l_i)}\right)}{\log(1-t_{\tau-1}}) + (1-\lambda)\cdot\epsilon^e_{i,\tau-1} \tag{A11}$$

$$= \lambda \cdot \frac{\log\left(\frac{(1-f(t))(t_{\tau-1}y_i + T_{\tau-1})}{(1-f(t))t_{\tau-1}(a\cdot\bar{y} + (1-a)l_i)}\right)}{\log(1-t_{\tau-1})} + (1-\lambda)\cdot\epsilon^e_{i,\tau-1} \tag{A12}$$

$$= \lambda \cdot \frac{\log\left(\frac{t_{\tau-1}y_i + T_{\tau-1}}{t_{\tau-1}(a\cdot\bar{y} + (1-a)l_i)}\right)}{\log(1-t_{\tau-1})} + (1-\lambda)\cdot\epsilon^e_{i,\tau-1} \tag{A13}$$

$$= \epsilon^e_{i,\tau}. \tag{A14}$$

The updating mechanism is thus also independent of the rescaling by $(1-f(t))$ and $\epsilon^e_{i,\tau}(f) = \epsilon^e_{i,\tau}$. Finally, consider the updating rule for local perceptions:

$$l^e_{i,\tau}(f) = \lambda \cdot \frac{a\bar{y}(f) - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i(f) + T_{\tau-1}(f))}{t_{\tau-1}}}{a-1} + (1-\lambda)l^e_{i,\tau-1} \tag{A15}$$

Substituting the expression for the actual transfer $T_{\tau-1}(f) = (1-f(t))T_{\tau-1}$ as well as $y_i(f)$ and $\bar{y}(f)$ into the expression yields

$$l^e_{i,\tau}(f) = \lambda \cdot \frac{a(1-f(t))\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}(1-f(t))y_i + (1-f(t))T_{\tau-1})}{t_{\tau-1}}}{a-1} + (1-\lambda)l^e_{i,\tau-1} \tag{A16}$$

$$= \lambda(1-f(t))\frac{a\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i + T_{\tau-1})}{t_{\tau-1}}}{a-1} + (1-\lambda)l^e_{i,\tau-1}. \tag{A17}$$

For initialization, we set $l^e_{i,0}(f) = l_i(f) = (1-f(t))l_i = (1-f(t))l^e_{i,0}$. It follows that

$$l^e_{i,1}(f) = \lambda(1-f(t))\frac{a\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i + T_{\tau-1})}{t_{\tau-1}}}{a-1} + (1-\lambda)(1-f(t))l^e_{i,0} \tag{A18}$$

$$= (1 - f(t))l^e_{i,1} \tag{A19}$$

Going one period forward, we find that

$$l^e_{i,2}(f) = \lambda(1 - f(t))\frac{a\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i+T_{\tau-1})}{t_{\tau-1}}}{a-1} + l^e_{i,1}(f) \tag{A20}$$

$$= \lambda(1 - f(t))\frac{a\bar{y} - \frac{(1-t_{\tau-1})^{-\epsilon}(t_{\tau-1}y_i+T_{\tau-1})}{t_{\tau-1}}}{a-1} + (1-\lambda)(1 - f(t))l^e_{i,1} \tag{A21}$$

$$= (1 - f(t))l^e_{i,2}. \tag{A22}$$

By the same argument, it follows that $l^e_{i,\tau}(f) = (1 - f(t))l^e_{i,\tau}$ for all $\tau = 0, 1, \ldots$, i.e., local perceptions are rescaled by $(1 - f(t))$. Substituting this expression into the voting rule yields

$$V_{i,\tau}(f) = \begin{cases} 1\text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y}(f) + (1-a) \cdot l^e_{i,\tau}(f)] > y_i(f) \\ 0\text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y}(f) + (1-a) \cdot l^e_{i,\tau}(f)] = y_i(f) \\ -1\text{if } (1-t)^\epsilon \cdot [a \cdot \bar{y}(f) + (1-a) \cdot l^e_{i,\tau}(f)] < y_i(f), \end{cases} \tag{A23}$$

which implies that

$$V_{i,\tau}(f) = \begin{cases} 1\text{if } (1-t)^\epsilon \cdot (1-f(t)) \cdot [a \cdot \bar{y} + (1-a) \cdot l^e_{i,\tau}] > (1-f(t))y_i \\ 0\text{if } (1-t)^\epsilon \cdot (1-f(t)) \cdot [a \cdot \bar{y} + (1-a) \cdot l^e_{i,\tau}] = (1-f(t))y_i \\ -1\text{if } (1-t)^\epsilon \cdot (1-f(t)) \cdot [a \cdot \bar{y} + (1-a) \cdot l^e_{i,\tau}] < (1-f(t))y_i, \end{cases} \tag{A24}$$

and, consequently, that $V_{i,\tau} = V_{i,\tau}(f)$, as the factor $(1 - f(t))$ cancels out on both sides of the (in)equalities. Homogeneous tax avoidance thus does not affect voting behavior for updating beliefs about the local income. The results we provide in the main text for full tax enforcement, therefore, hold for the general case with successful tax avoidance by agents. Note that this result is crucially dependent on the assumption that tax avoidance is homogeneous and does, in particular, not depend on the level of gross income y_i. If that is not the case, the factor $(1 - f(t, y))$ will no longer cancel out on both sides of the inequalities for the voting decision and might alter agents' votes. An interesting extension of the model might thus be to integrate income-dependence into the tax avoidance decision or even to let the avoidance decision be affected by the decisions of an agent's link-neighbors.

Appendix C

Dynamic voting behavior for varying weight of the global signal

See Fig. 13.

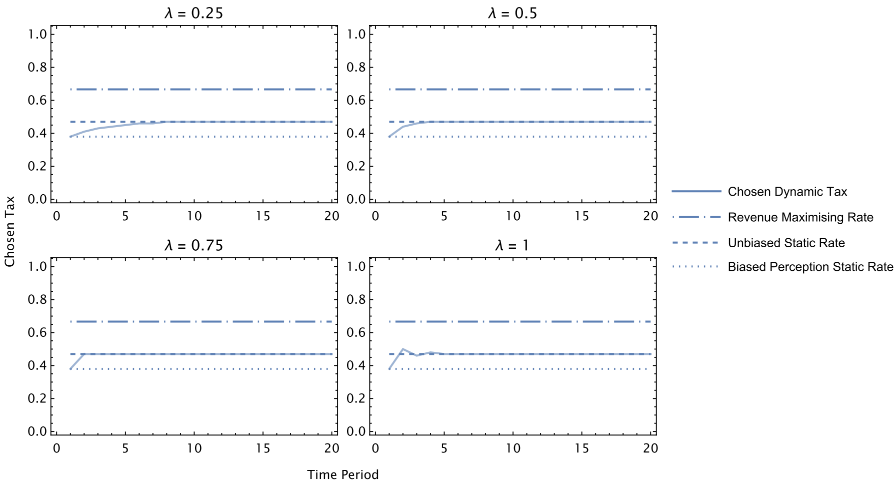

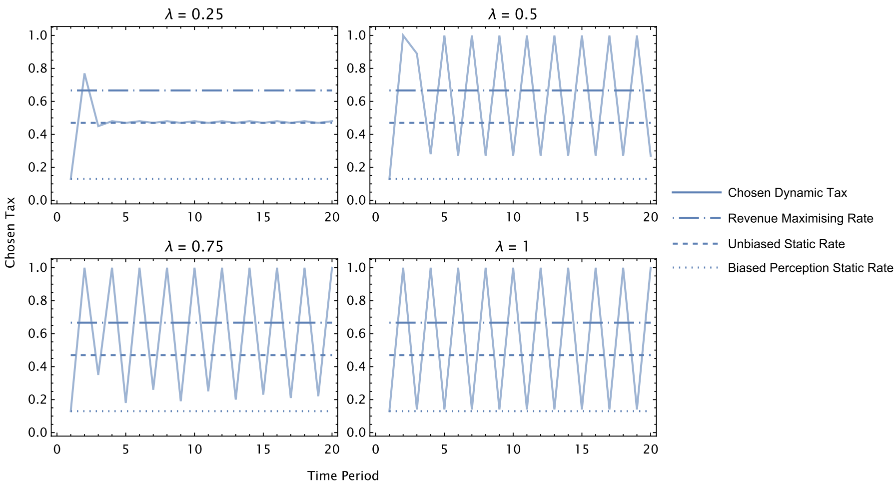

Fig. 13 The figure shows the chosen tax rates for a dynamic updating process for beliefs about ϵ with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$. All simulations are conducted for $\rho = 8$ and a true $\epsilon = 0.5$. The upper four panels correspond to the case of "high" $a = 0.75$, 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 behavior for the low a is consistently 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 and Fig. 15

Appendix D

Dynamic voting behavior for taxation inefficiency with bounded beliefs

See Figs. 14, 15.

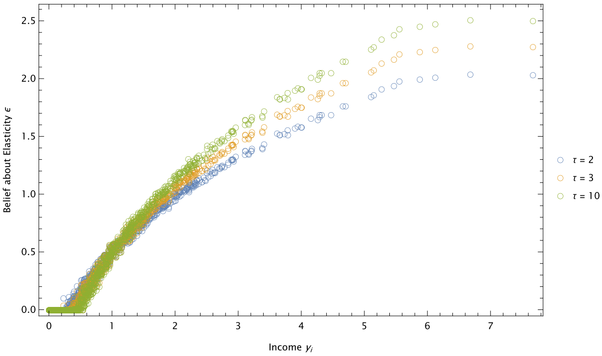

Fig. 14 The figure shows agents' beliefs about the elasticity of taxable income ϵ against their pre-tax incomes y_i for $\rho = 8$, $a = 0.5$, $\lambda = 0.25$ and the true $\epsilon = 0.5$. Beliefs for poor agents converge to zero quickly, while the beliefs of richer agents are unaffected compared to the benchmark without bounds in Fig. 9

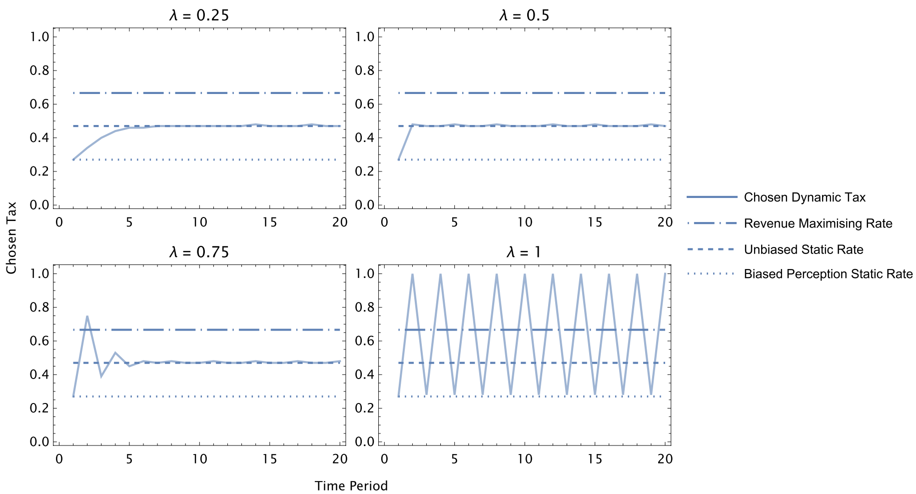

Fig. 15 The figure shows the chosen tax rates for the dynamic updating process for ϵ with different error-adjustment parameters $\lambda \in \{0.25; 0.5; 0.75; 1\}$ with a lower bound for $\epsilon_i \geq 0$. All simulations are conducted for $a = 0.5$, $\rho = 8$ and a true $\epsilon = 0.5$. As is immediately visible from the figure compared to the scenario without such a lower bound in Fig. 10, the lower bound on ϵ_i has no effect on the chosen tax rates

Acknowledgements We thank Bettina Gregg, Torsten Heinrich, Leonhard Ipsen, Johannes Marx, Nils Rochowicz, and Moritz Schulz as well as the participants in the Research Seminar on Political Theory of the University of Bamberg, the Interdisciplinary Workshop on Opinion Dynamics and Collective Decision 2023 in Bremen, the 34th Conference of the European Association for Evolutionary Political Economy, the 2024 Annual Meeting of the Keynes Society in Bamberg, the workshop of the AK Political Theory 2024 in Bamberg, and the Fourth Pluralumn* Workshop in Bamberg for their valuable comments. We also thank Daniel Hilbinger, Linus Glaser and Marie Schwab for their able research assistance. An earlier version of our research was presented at the 2nd Workshop in Agent-Based Modeling & Policy-Making in Saarbrücken, where it greatly benefited from the discussion and valuable points raised by five anonymous reviewers.

Funding Mayerhoffer's and Schulz's contributions were funded by the German Federal Ministry of Research, Technology and Space (grant: DATI Pilot, e-lernen). Mayerhoffer's contribution was co-funded by the Data Science Centre of the University of Amsterdam. Research for the paper was done during Mayerhoffer's fully-funded stay as International Diversity Guest Professor at the University of Bamberg.

References

Aguiar-Conraria, L., P. C. Magalhães, and M. J. Soares. 2012. Cycles in politics: wavelet analysis of political time series. *American Journal of Political Science* 56 (2): 500–518.

Aguiar-Conraria, L., P. C. Magalhaes, and M. J. Soares. 2013. The nationalization of electoral cycles in the united states: a wavelet analysis. *Public Choice* 156:387–408.

Alesina, A., and G. M. Angeletos. 2005. Fairness and redistribution. *American Economic Review* 95 (4): 960–980.

Alstadsæter, A., N. Johannesen, and G. Zucman. 2019. Tax evasion and inequality. *American Economic Review* 109 (6): 2073–2103.

Anufriev, M., and C. Hommes. 2012. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics* 4 (4): 35–64.

Arndt, H. L. R. 2020. Varieties of affluence: How political attitudes of the rich are shaped by income or wealth. *European Sociological Review* 36 (1): 136–158.

Atria, J. 2023. No taxation without efficiency? Elite perceptions of redistribution and progressivity in Chile. *Journal of Latin American Studies* 55 (1): 103–128.

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

Barnes, L. 2022. Taxing the rich: public preferences and public understanding. *Journal of European Public Policy* 29 (5): 787–804.

Bartels, L.M. 2009. Economic inequality and political representation. *The unsustainable American state*: 167–196 .

Bastani, S., and J. Lundberg. 2017. Political preferences for redistribution in Sweden. *Journal of Economic Inequality* 15 (4): 345–367.

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

Borge, L. E., and J. Rattsø. 2004. Income distribution and tax structure: Empirical test of the Meltzer-Richard hypothesis. *European Economic Review* 48 (4): 805–826.

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

Brooks, C., and D. Brady. 1999. Income, economic voting, and long-term political change in the US, 1952–1996. *Social Forces* 77 (4): 1339–1374.

Caplan, B. 2001. Rational irrationality and the microfoundations of political failure. *Public Choice* 107 (3–4): 311–331.

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 D. Fichera. 2020. Tax evasion and tax morale: A social network analysis. *European Journal of Political Economy* 65 : 101922.

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

Dimick, M., D. Rueda, and D. Stegmueller. 2018. Models of other-regarding preferences, inequality, and redistribution. *Annual Review of Political Science* 21 (1): 441–460.

Drăgulescu, A., and V. M. Yakovenko. 2001. Evidence for the exponential distribution of income in the USA. *The European Physical Journal B* 20 (4): 585–589.
Dubois, E. 2016. Political business cycles 40 years after Nordhaus. *Public Choice* 166:235–259.
Enke, B. 2020. What you see is all there is. *Quarterly Journal of Economics* 135 (3): 1363–1398.
Evans, G., and S. Honkapohja. 2001. Expectations. In *economics of, In International Encyclopedia of the Social & Behavioral Sciences*, ed. N. J. Smelser and P. B. Baltes, 5060–5067. Oxford: Pergamon.
Foley, D. K. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7:45–98.
Franko, W. W., and A. C. Livingston. 2022. Economic segregation and public support for redistribution. *Social Science Journal* 59 (3): 375–393.
Galesic, M., H. Olsson, and J. Rieskamp. 2012. Social sampling explains apparent biases in judgments of social environments. *Psychological Science* 23 (12): 1515–1523.
Georgiadis, A., and A. Manning. 2012. Spend it like Beckham? Inequality and redistribution in the UK, 1983–2004. *Public Choice* 151:537–563.
Grüne-Yanoff, T., and P. Verreault-Julien. 2021. How-possibly explanations in economics: anything goes? *Journal of Economic Methodology* 28 (1): 114–123.
Guerrero, A. 2021. The epistemic pathologies of elections and the epistemic promise of lottocracy, *Political Epistemology*, 156–179. Oxford University Press.
Hasenfeld, Y., and J. A. Rafferty. 1989. The determinants of public attitudes toward the welfare state. *Social Forces* 67 (4): 1027–1048.
Hommes, C. 2011. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control* 35 (1): 1–24.
Hommes, C., D. Massaro, and I. Salle. 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 K. Gurney. 2008. Network 'small-world-ness': a quantitative method for determining canonical network equivalence. *PloS One* 3 (4) : e0002051.
Iacono, R., and M. Ranaldi. 2021. The nexus between perceptions of inequality and preferences for redistribution. *The Journal of Economic Inequality* 19 (1): 97–114.
Jachimowicz, J.M., S. Davidai, D. Goya-Tocchetto, B. Szaszi, M.V. Day, S.J. Tepper, L.T. Phillips, M.U. Mirza, N. Ordabayeva, and O.P. Hauser. 2022. Inequality in researchers' minds: Four guiding questions for studying subjective perceptions of economic inequality. *Journal of Economic Surveys* 37(5):1534–1561
Knell, M., and H. Stix. 2020. Perceptions of inequality. *European Journal of Political Economy* 65 : 101927.
Landgren, E., J. L. Juul, and S. H. Strogatz. 2021. How a minority can win: Unrepresentative outcomes in a simple model of voter turnout. *Physical Review E* 104 (5) : 054307.
Larcinese, V. 2007. Voting over redistribution and the size of the welfare state: the role of turnout. *Political Studies* 55 (3): 568–585.
Lebo, M. J., and H. Norpoth. 2007. The pm and the pendulum: dynamic forecasting of british elections. *British Journal of Political Science* 37 (1): 71–87.
Lindert, P. H. 1996. What limits social spending? *Explorations in Economic History* 33 (1): 1–34.
Mac Carron, P., K. Kaski, and R. Dunbar. 2016. Calling Dunbar's numbers. *Social Networks* 47:151–155.
Mayerhoffer, D. M., and J. Schulz. 2022. Perception and privilege. *Applied Network Science* 7 (1): 1–25.
Mayerhoffer, D.M., and J. Schulz 2023. Redistribution, social segregation and voting information. In G. Sileno, N. Lettieri, and C. Becker (Eds.), *Proceedings of the 2nd Workshop on Agent-based Modeling and Policy-Making (AMPM 2022)*. CEUR Workshop Proceedings.
Mayerhoffer, D.M. and J. Schulz. 2025. Inequality perceptions (model code). URL: github.com/mayerhoffer/Inequality-Perception.
McPherson, M., L. Smith-Lovin, and J. M. Cook. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27 (1): 415–444.
Meltzer, A. H., and S. F. Richard. 1981. A rational theory of the size of government. *Journal of Political Economy* 89 (5): 914–927.
Meltzer, A. H., and S. F. Richard. 1983. Tests of a rational theory of the size of government. *Public Choice* 41 (3): 403–418.
Merrill, S., B. Grofman, and T. L. Brunell. 2008. Cycles in American national electoral politics, 1854–2006: statistical evidence and an explanatory model. *American Political Science Review* 102 (1): 1–17.
Milanovic, B. 2000. The median-voter hypothesis, income inequality, and income redistribution: an empirical test with the required data. *European Journal of Political Economy* 16 (3): 367–410.
Minkoff, S. L., and J. Lyons. 2019. Living with inequality: Neighborhood income diversity and perceptions of the income gap. *American Politics Research* 47 (2): 329–361.
Mäki, U. 2009. MISSing the world. Models as isolations and credible surrogate systems. *Erkenntnis* 70:29–43.
Nordhaus, W. D. 1975. The political business cycle. *Review of Economic Studies* 42 (2): 169–190.

Nordhaus, W. D., A. Alesina, and C. L. Schultze. 1989. Alternative approaches to the political business cycle. *Brookings Papers on Economic Activity* 1989 (2): 1–68.

Norpoth, H. 1995. Is Clinton doomed? An early forecast for 1996. *PS: Political Science & Politics 28*(2): 201–207 .

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 A. Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64 (6): 1263–1297.

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

Sabet, N. 2023. Turning out for redistribution: the effect of voter turnout on top marginal tax rates. *Public Choice* 194 (3–4): 347–367.

Sands, M. L. 2017. Exposure to inequality affects support for redistribution. *Proceedings of the National Academy of Sciences* 114 (4): 663–668.

Sands, M. L., and D. de Kadt. 2020. Local exposure to inequality raises support of people of low wealth for taxing the wealthy. *Nature* 586 (7828): 257–261.

Scervini, F. 2012. Empirics of the median voter: democracy, redistribution and the role of the middle class. *Journal of Economic Inequality* 10:529–550.

Schokkaert, E., and T. Truyts. 2017. Preferences for redistribution and social structure. *Social Choice and Welfare* 49:545–576.

Schulz, J., and D. Mayerhoffer. 2024. Social sampling, self-anchoring and redistribution. *Economics Letters* 244 : 111942.

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

Schulz, J., and D. M. Mayerhoffer. 2023. A network approach to expenditure cascades. *Review of Behavioral Economics* 10 (3): 229–262.

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

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

Sumaktoyo, N. G., C. Breunig, and W. Gaissmaier. 2022. Social sampling shapes preferences for redistribution: Evidence from a national survey experiment. *Journal of Experimental Social Psychology* 101 : 104341.

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

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

Williamson, V. 2019. Public ignorance or elitist jargon? Reconsidering Americans' overestimates of government waste and foreign aid. *American Politics Research* 47 (1): 152–173.

Windrum, P., G. Fagiolo, and A. Moneta. 2007. Empirical validation of agent-based models: Alternatives and prospects. *Journal of Artificial Societies and Social Simulation* 10 (2): 8.

Windsteiger, L. 2022. The redistributive consequences of segregation and misperceptions. *European Economic Review* 144 : 104073.

Xu, P., and J. C. Garand. 2010. Economic context and Americans' perceptions of income inequality. *Social Science Quarterly* 91 (5): 1220–1241.