Applied Network Science

Perception and privilege

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

*Correspondence:
daniel.mayerhoffer@uni-
bamberg.de

[1] Institute for Political
Science, University
of Bamberg, Bamberg,
Germany
Full list of author information
is available at the end of the
article

Abstract

Inequality perceptions differ along racial and gendered lines. To explain these dispari-
ties, we propose an agent-based model of localised perceptions of the gender and
racial wage gap in networks. We show that the combination of homophilic graph
formation and estimation based on locally limited knowledge can replicate both the
underestimation of the gender or racial wage gap that empirical studies find and the
well-documented fact that the underprivileged perceive the wage gap to be higher on
average and with less bias. Similarly, we demonstrate that the underprivileged perceive
overall inequality to be higher on average. In contrast to this qualitative replication, we
also show that the effect of homophilic graph formation is quantitatively too strong to
account for the empirically observed effect sizes within a recent Israeli sample on per-
ceived gender wage gaps. As a parsimonious extension, we let agents estimate using
a composite signal based on local and global information. Our calibration suggests
that women place much more weight on the (correct) global signal than men, in line
with psychological evidence that people adversely affected by group-based inequities
pay more attention to global information about the issue. Our findings suggest that
(educational) interventions about the global state of gender equality are much more
likely to succeed than information treatments about overall inequality and that these
interventions should target the privileged.

Keywords: Inequality, Homophily, Disassortativity, Diversity, Gender, Race, Wage gap,
Agent-based model

Introduction

Perceptions matter. There is now widespread agreement that it is *perceived* inequi-
ties, not necessarily actual ones behind redistributive preferences and policy decisions
(Hauser and Norton 2017; Choi 2019). Apart from biased perceptions of empirical
wealth and income distributions, the literature increasingly also documents mispercep-
tions of the gender and racial inequities (Kraus et al. 2017; Malul 2021), relating, in par-
ticular, to wage gaps between groups of the privileged and underprivileged. This finding
carries a vital policy implication: If the general public overestimates the achieved pro-
gress in both racial and gender equality, the issue might never surface in public debate
(Ceccardi 2021). There also exists considerable theoretical and empirical evidence that
individual perceptions of unfairness trigger wage negotiations and, consequently, wage
growth (Akerlof and Yellen 1990; Pfeifer and Stephan 2019). Misperceptions about the

wage gap might thus contribute to its persistence documented in Akchurin and Lee (2013)—and addressing this persistence would require addressing its biased perceptions.

Despite their relevance for both market outcomes and policy, there is no formal model of belief formation in this regard to the best of our knowledge. There are two potential channels that may (jointly or independently) lead to the biased individual perceptions: firstly, biased information processing of a correct signal, i.e., the true global (population-wide) wage gap; secondly, incorrect individual signals that come from a skewed sample. Such a sample might fail to accurately represent the population, e.g., due to social networks segregated by income (Kraus et al. 2017).

We propose a simple network model based on homophilic attachment initially employed for perceptions of overall inequality by Schulz et al. (2022) and find that the second channel alone is sufficient to explain an underestimation of the wage gap, even given fully rational, unbiased individual information processing. In fact, model agents who rely only on the biased information from their individual sample underestimate the wage gap more severely than empirical studies find. Thus, we combine the biased local signal with the correct wage gap as global signal.

By calibrating our model with data from a recent sample in Malul (2021), we find that both the local and the global channels matter but that their relative importance differs between the privileged and underprivileged. Our theoretical findings imply that informational treatments appear to have a much more significant effect when targeted on the privileged.

By contrast, our results suggest that purely local perceptions are sufficient to generate the empirically observed underestimation of overall inequality. Here, we find that the underprivileged unambiguously perceive higher inequality and are more accurate on average in their assessment. This insight offers a novel explanation for why women and ethnic minorities consistently favour more redistribution. Rather than assuming that they have different value orientations or want the welfare state to redistribute more out of self-interest, we argue that they simply have a clearer picture of the problem and perceive inequality to be higher.

The remainder of this paper is organised as follows: Section "Related literature" discusses related empirical findings that our model aims to replicate. We introduce the model in section "Model". Section "Results" introduces our basic results both on the mean perceptions for the privileged and underprivileged class as well as the highly non-monotonic relationship between income, neighbourhood disassortativity and perceptions, as well as our calibration exercise with Israeli data from Malul (2021). Section "Discussion" spotlights policy implications of our results and discusses avenues for further theoretical and empirical work.

Related literature

This paper considers perceptions of the wage gap between the privileged and underprivileged as emerging from endogenously evolving reference groups. The yet relatively scarce literature on such perceived wage gaps has mainly focused on gender and racial wage gaps. For the gender wage gap, the extant literature suggests that the population appears to underestimate the actual extent of gender-based inequality (Malul 2021) but that this effect is primarily driven by male bias, with women having much higher and

more accurate perceptions on average (Hampton and Heywood 1993; Foley et al. 2002; García-González et al. 2019; Malul 2021). In line with these findings, women exhibit a stronger tendency to perceive both the overall level of inequality (Andreoni and Vesterlund 2001; Alesina and Giuliano 2011; Bavetta et al. 2019) and gender-based inequalities (Pfeifer and Stephan 2019) to be unjustified.[1] Regarding perception formation about the gender wage gap, the literature suggests two predominant channels: Individuals form their perceptions from lived experience in (potentially) gender-diverse situations (Auspurg et al. 2017) but are also generally aware of the global wage gap, e.g., through national media (Williams et al. 2010; Furnham and Wilson 2011).

The findings on perceptions of racial inequality are largely in line with those for the gender wage gap. The general population underestimates the extent of racial inequality across various dimensions, including the wage gap (Kraus et al. 2017, 2019; Alesina et al. 2021; Davidai and Walker 2021). These studies also tend to find that Black people who are adversely affected by the wage gap tend to perceive higher wage gaps and are, generally, more accurate in its assessment (Kraus et al. 2017, 2019; Davidai and Walker 2021). Kraus et al. (2017) also find that the richest White people tend to underestimate the racial gap most strongly. Intriguingly, they speculate and provide some evidence that this might be related to the diversity in social networks that is particularly low for the wealthiest part of the population. Carter et al. (2019) corroborate this hypothesis and show that diversity in students' friendship networks empirically exhibits a positive association with perceptions of racial injustice. Apart from this local channel, the efficacy of informational treatments reporting aggregates also suggests the relevance of global averages on individual perceptions (Haaland and Roth 2021).

Homophily in social contacts may be the product of chance or a by-product of individual behaviour with different aims—or individuals may be consciously homophilic in their tie formation, e.g., to reduce strategic uncertainty (Kets and Sandroni 2019). Empirically, there is explicit income homophily surfacing in, e.g. mobile phone communication (Xu et al. 2021; Leo et al. 2016; Fixman et al. 2016). Moreover, willing and accidental homophily in other dimensions can represented as income homophily because income proximity is a proxy for proximity in other dimensions: Empirical studies show factors like residential area (Hu and Liang 2022; Harting and Radi 2020), education (Smith et al. 2014; Leo et al. 2016), lifestyle (Virtanen et al. 2007), or even health (Krieger 1992) to correlate with income; and there is obviously also a correlation between income and ethnicity (Chandra 2000), i.e. one possible dimension to observe privilege on with regard to our model. Of course, the named factors are also correlated with each other and we do not aim to argue for a potential direction of causality. Instead, we utilise the implication of the empirically well established correlation: Even if people do not willingly choose their social contacts based on income proximity, their actual choices amount to social contacts as if selected homophilic in income.

To validate our model that we introduce in section "Model", we thus consider two stylised facts extracted from the empirical literature: (i) The general population appears to

---

[1] Note that the effect of gender on perceptions of fairness reverses in Pfeifer and Stephan (2019) when controls for the hourly wage are introduced. Since our model-output consists of the unconditional effects without controlling for income, the unconditional effect in Pfeifer and Stephan (2019) is the relevant benchmark, though.

underestimate the extent of gender or racial inequality and (ii) this underestimation is mainly driven by the bias of the privileged agents. As for generating mechanisms, both local knowledge in the form of lived experience and global signals, e.g. through the national media, appear to be necessary.

While perceptions thus appear to be detached from actual gender and racial inequities, we still need to consider the actual levels of inequality as a benchmark. Gender and racial wage gaps were generally between 10 % and 40% in the last decade for industrialised countries which is also the range we will consider (Kunze 2018; Stelzner and Bahn 2020). Moreover, it is now a well-established stylised fact that empirical wage distributions follow an exponential law which has been demonstrated in numerous studies for a vast range of countries and sample periods, i.e., for Australia (Banerjee et al. 2006), Romania (Derzsy et al. 2012; Oancea et al. 2017), the European Union as a whole (Jagielski and Kutner 2013), Japan (Nirei and Souma 2007), the US (Nirei and Souma 2007; Dragulescu and Yakovenko 2001a, b; Banerjee and Yakovenko 2010; Dos Santos 2017; Schneider 2015; Shaikh 2017; Shaikh and Jacobo 2020; Silva and Yakovenko 2004) as well as the UK (Dragulescu and Yakovenko 2001b; Banerjee and Yakovenko 2010; Dos Santos 2017; Schneider 2015; Shaikh 2017; Shaikh and Jacobo 2020; Silva and Yakovenko 2004). Most notably, Tao et al. (2019) confirm for their whole sample of 67 countries that the exponential law provides an excellent fit for empirical wage distributions.[2] For our purposes, though, not only the overall distribution is relevant but also the wage distribution of privileged and underprivileged wage-earners considered separately. In a foundational contribution, Shaikh et al. (2014) indeed also demonstrated that the wage distribution of men and women, as well as the ones of Whites, Hispanics and African-American considered separately, are well approximated by this exponential distribution for US data. While the wage distributions of the underprivileged class thus have lower average income, they empirically follow the same functional form as the wage distributions of the privileged and the population as a whole.

After an in-depth validation of our model, we directly apply the general model structure to a long-standing policy issue: Attitudes towards redistribution. A large literature now documents that women (Linos and West 2003; Quadagno and Blekesaune 2003; Cusack et al. 2006; Alesina and Giuliano 2011) and ethnic minorities (Alesina and Giuliano 2011; Kinder and Winter 2001; Morgan and Kelly 2017) are generally much more supportive of welfare state measures and redistributive policies. The main hypotheses to explain these gender and racial differentials are based on different value orientations or on the fact that the precarious working and living conditions by women and ethnic minorities let those groups support redistributive measures out of self-interest (Quadagno and Blekesaune 2003). Similarly, Linos and West (2003) suggest that women tend to support welfare state measures because they are more likely to work within the public sector, therefore directly benefitting from the expansion of the welfare state. We provide a complementary explanation: Our model unambiguously implies that the underprivileged perceive higher levels of inequality and are less biased in their

---

[2] The upper tail of 1–5% empirical income distributions appears to follow a power-law or Pareto distribution (Silva and Yakovenko 2004). The income concentrated there is primarily *capital income*, though, while we are concerned with *wages* and labour income. Henceforth, we exclude this small minority from our consideration.

Fig. 1 Overview over the simulation processes
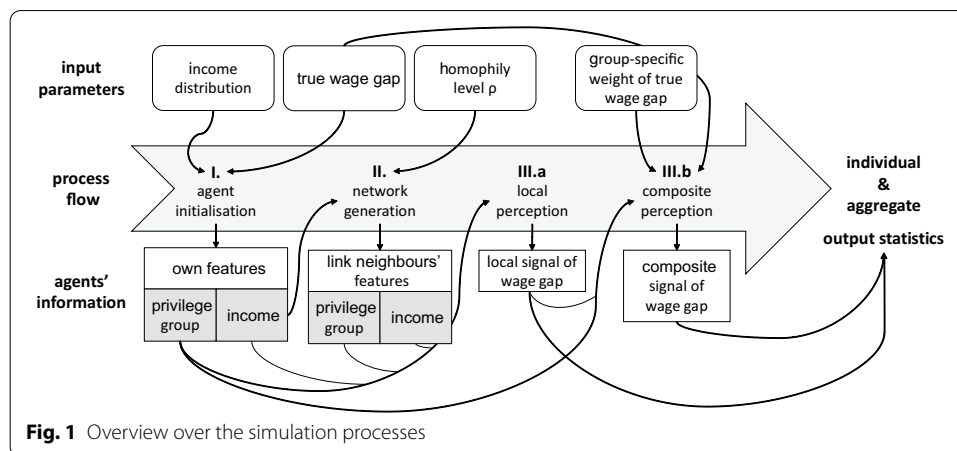
perception. Thus, they might favour redistribution simply because they have a much clearer picture of the existing inequalities and encounter them more frequently or saliently in their daily lives.

Model

The present agent-based computational model aims to study the impact of social tie formation that is homophilic in wage on individual perceptions of a wage gap and wage inequality in general. The link formation is important because model agents use their ego network as a sample to form their perceptions. The model and analyses here extend (Mayerhoffer and Schulz 2022) who in turn build on a linking procedure introduced by Schulz et al. (2022) to study perceptions of inequality in general. It comprises of three distinct phases that take place once per simulation run and in sequential order: (I) Agent initialisation and group-specific income allocation, (II) network generation through agents' homophilic linkage, and (III) individual wage gap perception and network evaluation, which comprises of (III.a) comparing wages between groups within one's neighbourhood and (III.b) processing of the correct population-wide wage gap as a global signal. Figure 1 gives an overview of these phases to guide the content-oriented descriptions below; for further technical details, consider the NetLogo implementation of the simulation model and the output data.

Agent initialisation

The simulated population consists of 1000 agents in the model that belong to either the privileged group/class (e.g., White people, men) or the underprivileged one (e.g., BiPoC, women). At initialisation, each agent draws their income from an exponential distribution with a mean of $\lambda = 1$. This pre wage gap income distribution is identical for all agents regardless of their privilege group. As discussed in section "Related literature", such a distribution normalises the empirically observed (pre-tax or market) wage distributions in various industrialised countries for the vast majority of individuals (Tao et al. 2019).

We use the empirical finding that the wage distribution of the total population as well as of the releant population subgroups is exponential for our model initialisation. The

wage distribution we use for initialisation can thus be considered as constituting a representative sample taken from empirical wage distributions. The distribution of underprivileged earners is downscaled by a wage gap $g \in [0, 1]$, in line with the empirical findings of Shaikh et al. (2014). Privileged agents thus always retain the income they drew while the underprivileged agents' income is reduced by a proportion $g \in [0, 1]$. The findings on the general population's income distribution (Tao et al. 2019) as well as our simulation results ensure that for any wage gap, the unconditional distribution of incomes irrespective of privilege also follows an exponential law. They allow us to construct a general model that applies equally to perceived gender and racial wage gaps based on privilege classes, since decomposition of the income distribution along both these dimensions delivers the same exponential distributional regularity.

Network generation

Each agent selects five other agents as link-neighbours. Therefore, the model does not impose links on agents according to a global rule, but, like in real-world networks, they *create* their links themselves. If an agent i picks agent j who had themselves picked i before that, the already existing link between the two agents remains untouched, but i does not pick another neighbour instead of j. Consequently, each agent has at least five link-neighbours (i.e., social contacts) but may have more. This number is empirically validated, as we intend to represent the closest layer of intense contacts (Zhou et al. 2005; Hamilton et al. 2007; Mac Carron et al. 2016)[3]. Moreover, we have carried out sensitivity analyses and found a higher number of links to have little impact on overall simulation results. Since this closest layer may consist of differently related people for different individuals (Karlsson et al. 2005), a link in our model can represent any social relationship: It simply indicates that an agent knows their link neighbour's income. Because links are undirected and equally weighted within the model, this knowledge about income is always mutual.

When selecting link neighbours, an agent does not care about the potential selectee's group belonging. However, agents do care about the potential drawee's income, as link-generation is homophilic in income. Homophily generally describes the tendency to link to similar individuals. This mechanism is important for information diffusion (Larson 2017). Moreover (Milli 2021) employs homophily in an agent-based model as influence factor of perception. Income constitutes one particularly relevant dimension (McPherson et al. 2001) of homophily and the only dimension we consider here. Namely, agent j's weight in agent i's draw is denoted by Ω_{ij} and determined as follows:

$$\Omega_{ij} = \frac{1}{\exp[\rho \, |Y_j - Y_i|]} \tag{1}$$

The relative weights in the draws are a function of the homophily strength and the respective income levels: $Y \in \mathbb{R}^+$ denotes the income of an agent. $\rho \in \mathbb{R}_0^+$ denotes the homophily strength in income selection, externally set, and identical for all agents. $\rho = 0$ represents a random graph, and for an increasing value of ρ, an agent becomes ever

---

[3] For a recent review on the extensive literature on 'Dunbar's number', cf. the first section of Mac Carron et al. (2016).

more likely to pick link-neighbours with incomes closer to their own. Put differently, for each agent, there is a likelihood ranking over all other agents based on income differences. Which agent occupies which rank and especially to which privilege class they belong depends on Y and thus on the wage gap g. ρ determines how strongly Ω reacts to those income differences.

The link function's exponential character ensures that those with the largest income differences become unlikely picks even at low to moderate homophily strengths. The choice of an exponential weighting function might seem arbitrary, but upon closer inspection, we find that translated into the probability of i choosing j, this weighting is equivalent to the discrete choice approach developed and popularised by Manski (1981). The homophily parameter $\rho \in [0, \infty)$ is then simply the intensity of the choice parameter. To translate weights into probabilities, we normalise by all weights for all agents, i.e.,
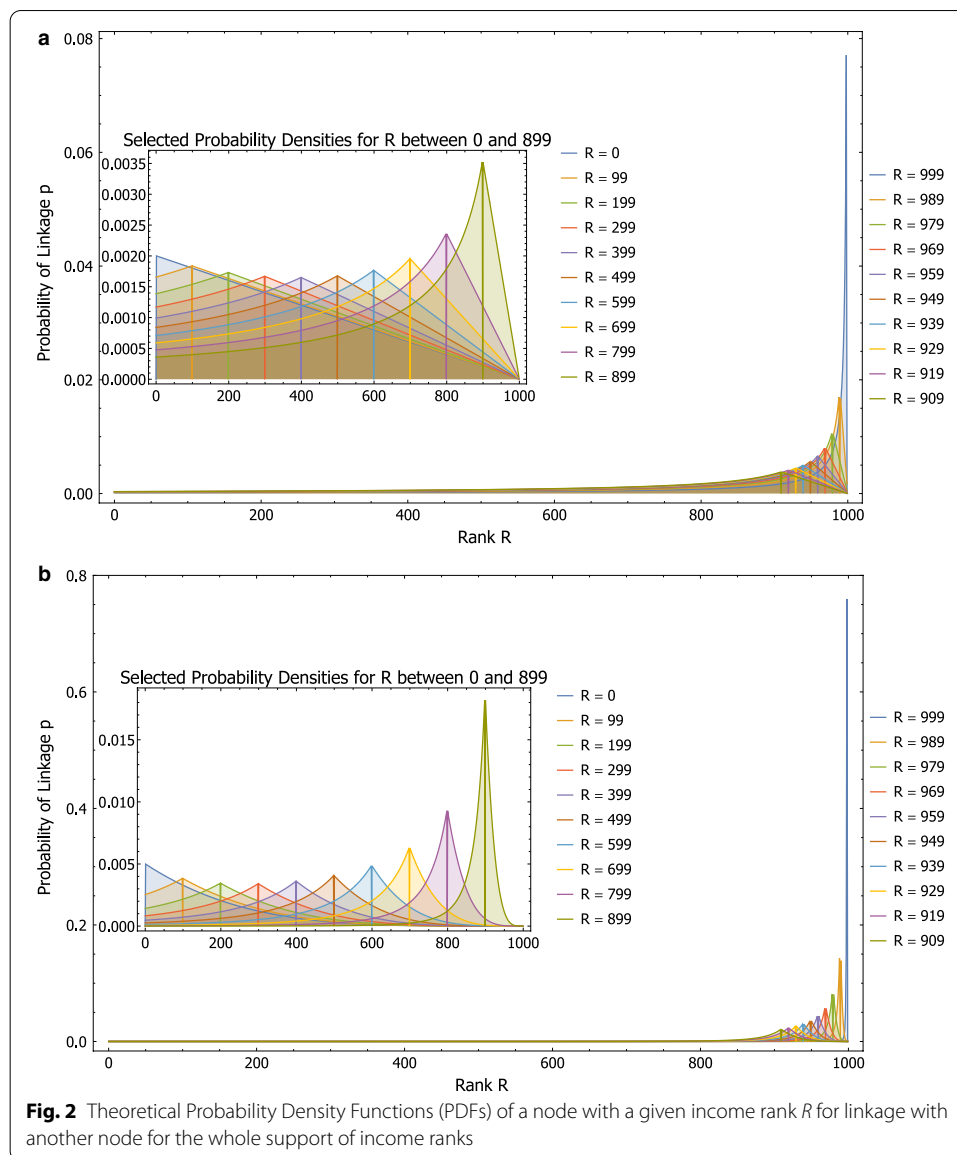
$$p_{ij} = \frac{\exp[-\rho \cdot |Y_j - Y_i|]}{\sum_{k \in M \setminus i} \exp[-\rho \cdot |Y_k - Y_i|]}, \tag{2}$$

with $M \setminus i$ as the set of all agents except i with size $N - 1$.[4] Considering probabilities rather than weights in (2) is intuitively interpretable, with $\rho = 0$ implying equiprobable picks with $p_{ij} = 1/(N - 1), \forall j \in M \setminus i$, and thus indeed a random graph, while $\rho \to \infty$ implies that p approaches unity for j with minimum income distance to the income of i and 0 for all other agents. Manski (1981) show that the discrete choice rule follows directly from random utility theory, i.e., agents maximising utility and utility functions being decomposable into an observable and unobservable component and both being uncorrelated. In our case, the observable component the agents minimise would be the income differences, with the unobservable part being all the attributes from which our agent in question would benefit due to their social connection. It follows that the linkage function in Eq. (1) is plausibly microfounded in a utility-maximising framework.[5] In this sense, our weighting function is a tractable reduced-form representation of the empirically established homophily. Falk and Knell (2004) demonstrate that endogenous reference standards like ours can also emerge from an optimisation framework with utility functions that incorporate plausible relative motives.

Figure 2 illustrates the linkage probabilities implied by the weighted draw based on the exponentially distributed income levels. The local maxima of individual linkage probability densities exhibit a bi-modal shape with peaks at the highest and lowest rank but are also heavily skewed to the left, i.e., agents with high incomes are most selective in their link picks. We understand selectivity as the tendency to select agents that are close in income rank as link-neighbours. For an extensive analysis of this linkage behaviour, see (Schulz et al. 2022). General selectivity increases with ρ, and for $g > 0$, selectivity in income may also mean selectivity in groups. Most notably, this implies that our model generates gender-based and racial homophily as a (from a

---

[4] Note that this is, strictly speaking, only the probability of the first choice of agent i, since we consider drawing without replacement. In particular, the weighting function does not account for the possibility that other agents already link to the agent in question, in contrast to our algorithm. Since the number of agents is rather large, the effect appears to cancel out in the aggregate, as is verified in (Schulz et al. 2022).

[5] The discrete choice approach has now become the workhorse formalisation within behavioural macroeconomics. Cf. e.g. Franke and Westerhoff (2017) for a recent survey.

Fig. 2 Theoretical Probability Density Functions (PDFs) of a node with a given income rank *R* for linkage with another node for the whole support of income ranks

modelling perspective desirable) byproduct of income homophily since income differences are partially determined by privilege class. Indeed, empirically, homophily operates along many different dimensions, including race and gender (McPherson et al. 2001).

Overall, an agent's choice of link neighbours depends on the homophily level *g* and the wage gap ρ. The resulting network is a member of the family of Random Geometric Graphs (Dall and Christensen 2002), which Talaga and Nowak (2019); Schulz et al. (2022) showed to reproduce core features of social networks efficiently, especially regarding their small-world character (Watts and Strogatz 1998). These include, in particular, a rather short average path lengths and high degrees of clustering. Specifically, the model combines the concept of homophily (Boguná et al. 2004) with presetting node degrees (Newman et al. 2001; Newman 2009). However, concerning our

application, we are able to simplify both approaches by pre-determination of only the global minimum degree, like in Preferential-Attachment networks, and consequently defining relative weights rather than absolute probabilities.

Individual perception

Like the linkage, any evaluation is done *by* the agents themselves given their individual knowledge. This knowledge consists of a local component defined by an agent's ego network and a global component, which is simply the true wage gap.

Comparison within one's neighbourhood

Agents know about their own income and also their social contacts' incomes. However, within the baseline specification, they do not possess knowledge about any other agent or structural features of the whole income distribution. Let U and P be the set of underprivileged and privileged agents and Θ_i the perception set of agent i consisting of $N_{U,i} \in \mathbb{N}_0$ underprivileged and $N_{P,i} \in \mathbb{N}_0$ privileged agents (including agent i themself).

Agent i then calculates the perceived local wage gap l_i as

$$l_i = \frac{\bar{Y}_P^i - \bar{Y}_U^i}{\frac{1}{2}(\bar{Y}_P^i + \bar{Y}_U^i)},$$

$$\text{with } \bar{Y}_P^i = (1/N_{P,i}) \cdot \sum_{j \in \{\Theta_i \cap P\}} Y_j, \bar{Y}_U^i = (1/N_{U,i}) \cdot \sum_{j \in \{\Theta_i \cap U\}} Y_j$$

and $N_{U,i} \neq 0$ as well as $N_{P,i} \neq 0$.

For either $N_{U,i} = 0$ or $N_{P,i} = 0$, we set $l_i = 0$, as in this case, agent i is unable to observe any wage gap. l_i is thus the percentage difference in perceived mean wages of the underprivileged and privileged agents within the perception set of i. The perceived local wage gap corresponds to the wage gap perception within any agent's ego network. We opt for the particular approximation above because it is symmetric (Törnqvist et al. 1985) and bounded in $[-2, 2]$ and thus more robust against outliers (Decker et al. 2014). Our results are not materially sensitive to the specific symmetric approximation of the growth rate and are very similar to the more common approximation by log-differences.[6]

Processing of a composite signal

Informed by our initial simulation results and the empirical finding that individuals are often aware of the wage gap through, e.g., national media, we consider the possibility that agents process a composite signal that combines their local perception and the global average wage gap. In this parsimonious extension, the composite perception p_i of individual i is a linear combination of local perceptions l_i and the (correct) global wage gap g by

$$p_i = (1 - w_i) \cdot l_i + w_i \cdot g. \tag{3}$$

The only free parameter is thus $w_i \in \mathbb{R}$, the weight i puts on the global signal. It is straightforward to solve for w_i as

---

[6] Results are available upon request.

$$w_i = \frac{p_i - l_i}{g - l_i},$$

(4)

with intuitive comparative statics. Whenever the composite perception p_i increases relative to l_i, the weight on the global signal increases, as local knowledge is less relevant for perceptions in these cases. Our dataset includes g and the mean perceptions for the underprivileged and privileged groups, i.e., $\bar{p}_U$ and $\bar{p}_P$. Of course, the aggregation function (4) in the form of a weighted average is not the only possible way to aggregate local and global signals. We chose this particular function form for its tractability and intuitive interpretability. Here, the respective weights w and $(1 - w)$ are immediately interpretable as percentage weights in contrast to more technically involved formulations.

Together with the simulation means for the local signal, $\bar{l}_U$ and $\bar{l}_P$, we are able to estimate the implied mean population weights $\bar{w}_U$ and $\bar{w}_P$ by replacing the individual parameters with the population means. In the Appendix, we show that using mean perceptions delivers exact results for the implied mean population weights, whenever the weights w within a subpopulation are uncorrelated to the local perceptions of the this group or, whenever the weights within each group of privileged and underprivileged individuals are homogenous, i.e., $\bar{w}_U$ is the same constant weight for all underprivileged agents and $\bar{w}_P$ is equal for all privileged agents. We consider deviations from this identification assumption in the discussion section below.

Perceiving overall inequality

We also consider the perception of *overall* inequality, i.e., the Gini coefficient rather than *intergroup* inequality as the wage gap between privilege groups. The local Gini perception mirrors the standard calculation of the Gini coefficient but is applied to the set of individuals i observes Θ_i and themself. The local Gini L_i is thus defined as

$$L_i = \frac{\Sigma_{j \in \Theta_i} \Sigma_{k \in \Theta_i} |Y_j - Y_k|}{2n^{-1} \Sigma_{j \in \Theta_i} Y_j}.$$

(5)

This local or perceived Gini coefficient L equals the ratio of the sum of all differences of incomes within the ego network of i (including agent i's income) to twice the mean income within this ego network (again calculated including agent i's income). Since the functional form of this calculation is equivalent to the standard Gini formula, bias only arises whenever the ego network is unrepresentative of the total population. This is the case for $\rho > 0$, with 'unrepresentativeness' increasing in ρ. For $\rho = 0$, perceptions are unbiased on average and irrespective of privilege class, as we also verify by simulation.

Validation strategy

To ensure internal validity of our model, i.e., detect any potential bugs, we obtained not only data on the wage gap but also general inequality perception and the network topology. This data is in line with what Schulz et al. (2022) find for their model without a wage gap, meaning that our model features the mechanisms that we intend. These mechanisms are to be understood as a (possible) minimal set of necessary assumptions (Grüne-Yanoff 2009) to explain the empirically observed underestimation of wage gaps. Hence, our model presents a how-possibly explanation looking for mechanisms that

could potentially cause the observed phenomenon and does not claim that these conditions are fulfilled in the real world. Nevertheless, there is a resemblance (Mäki 2009) between our model and the real world as the "model produces quantitative agreement with empirical macrostructures, as established through on-board statistical estimation routines" and also "quantitative agreement with empirical microstructures" (Barde and Van Der Hoog 2017). We use empirical micro-data to calibrate our model, namely an exponential income distribution and wage gaps that characterise industrialised countries as well as a linking function that is plausibly micro-founded within the discrete-choice approach. Furthermore, we compare the simulation output to available empirical macro-data. Therefore, the explanans can be true in the real world and the cause for the observed empirical fact: Our proposed mechanism fulfils the minimum conditions for a good epistemically possible how-possibly explanation formulated by Grüne-Yanoff and Verreault-Julien (2021). That our simulation results are in "qualitative agreement with empirical macrostructures" (Fagiolo et al. 2019, p. 771), namely successful replication of the stylised empirical facts, affirms the external validity of the model. Thus, the model is both technically verified as well as validated based on input and output measures. This combination of internal and external validity is a rare but desirable feature of simulation models (Gräbner 2018).

Results

We use $N = 1000$ agents for each simulation run that are equally split into a 'privileged' and 'underprivileged' class which thus contain 500 members each. 1000 agents ensure computability but is also in line with the sample sizes of empirical studies into wage distributions (Tao et al. 2019).[7] Both the income distribution for the privileged and underprivileged class are initialised by an identical exponential distribution that is characterised by a rate parameter $\lambda = 1$, i.e., mean income is also exactly equal to unity. All incomes within the distribution of the underprivileged class are then downscaled by the wage gap $(1 - g)$, as indicated by the empirical results in Shaikh et al. (2014). For our calibration by sample moments in subsection, "Empirical calibration and composite signal" we use the empirical estimates by Malul (2021), again assuming a total population $N = 1000$. Since our network generation mechanism is partly stochastic, our presented results are generally averages over 100 Monte Carlo (MC) runs each if not otherwise indicated.

Mean local perceptions

We first look at the mean local perceptions according to privilege and varying g and ρ in the empirically relevant ranges (without impact of a global signal). Recall that the scarce empirical literature suggests that (i) the wage gap is underestimated regardless of privilege but that (ii) the underprivileged have much more accurate perceptions than the privileged. Figure 3 summarises this first battery of simulation results.

In line with stylised fact (i), our simulated population overall tends to underestimate the wage gap considerably, regardless of privilege class, when we only consider

---

[7] Moreover, Schulz et al. (2022) carried out sensitivity analyses with higher N and found all results to be robust. They also showed analytically that all main model mechanisms are independent of sample size.

Fig. 3 Violin plots of perceived wage gap p_i relative to the true wage gap g for various g and ρ
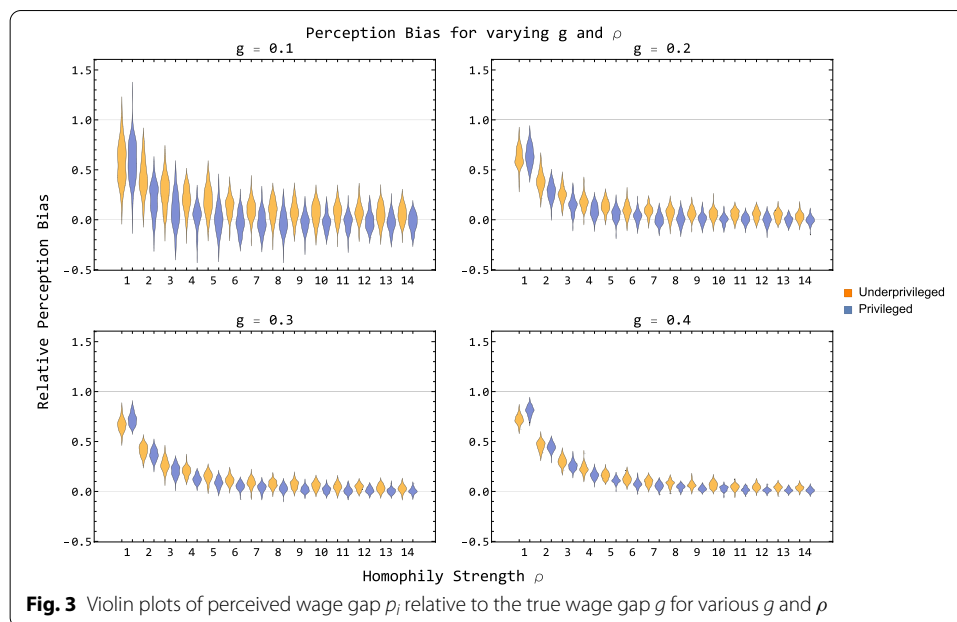
the local perceptions. We find that (almost) all violin plots are significantly below the gridline at unity that would indicate correct local perceptions. Of course, this outcome is entirely unsurprising. All agents estimate the wage gap based on the set of neighbours that have similar wage levels, since the graph formation is homophilic in income. It is thus improbable for any agent to observe incomes that strongly deviate from their own income in question, including most notably incomes from a different privilege class. Even though there is a unidirectional downwards bias, this bias affects the two privilege classes differently. Since selectivity is locally maximal at the lower and upper tail of the income distribution, as we show in Fig. 2, overall selectivity exhibits a U-shaped pattern. A wage gap $g > 0$ in combination with income homophily by $\rho > 0$ thus causes two counteracting effects within our model. First, since $g > 0$ moves the incomes of the richest underprivileged agents closer to the median income, it tends to decrease their selectivity, thus generally improving the accuracy of their wage gap estimation. Second, $g > 0$ also pushed the incomes of the poorest underprivileged agents away from the incomes of the poorest privileged agents and lets the poor underprivileged inhabit the lower tail of the income distribution by themselves. Thus, the neighbourhoods of the underprivileged poor are on average less diverse and their estimates of the wage gap are downwards biased. The direction of the composite effect depends on both g and ρ, as Fig. 3 shows. Generally, increasing ρ increases the relative strength of the first partial effect since the richest become much more selective here. Because our proposed mechanism is based entirely on income differences, increasing both g and ρ increases not only income segregation but also segregation based on privilege class. Conversely, the lower g and ρ, the more random the network generation will be and the higher the variability of simulation outcomes. As a stark illustration of this variability, note that the lowest $g = 0.1$ and $\rho = 1$ we consider is consistent with both an *overestimation* of the wage gap (simulation means above the

gridline at unity) and with *reversed* estimates where agents believe the privileged to earn less on average, i.e., mean estimates reaching below the gridline at zero. This effect is evident from the fact that variability in mean perceptions decreases both in ρ within each panel and in g across panels in Fig. 3.

We find that stylised fact (ii) is only possible in our model setup for homophily levels $\rho \geq 2$. This is in line with our findings for the perception of overall inequality, measured as the perceived Gini coefficient, that would imply $\rho \in [4, 14]$.[8] Note that homophily strengths in this range imply a tremendous downward bias in perceived wage gaps of at least 50%.
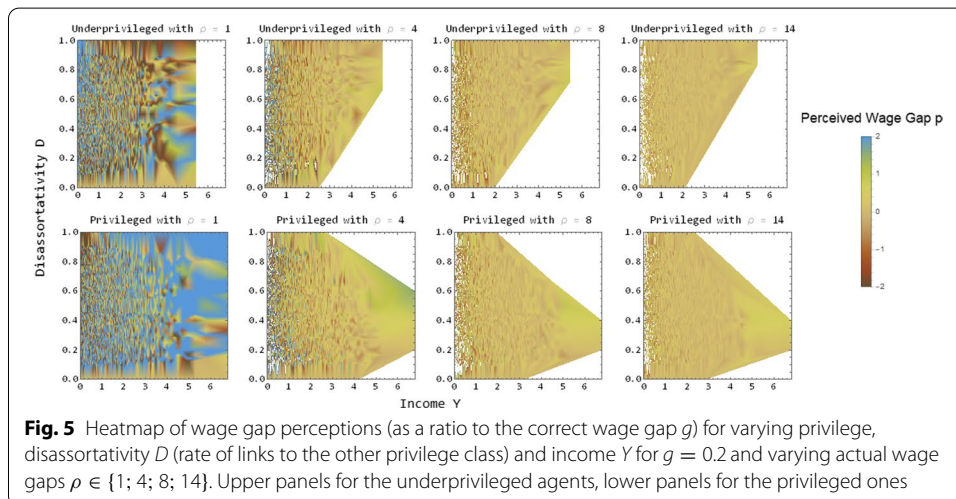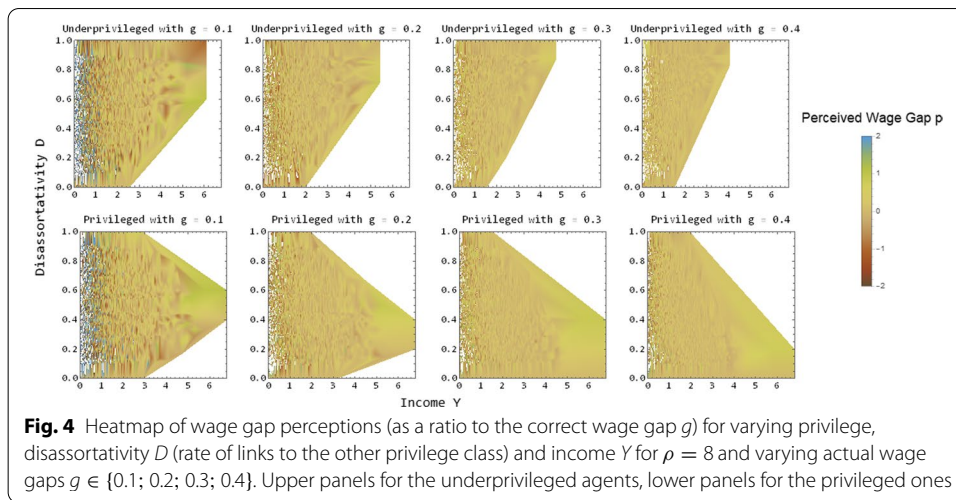
While we find that the two stylised facts (i) and (ii) can be *qualitatively* replicated for homophily strengths $\rho \geq 2$, our *quantitative* estimation results indicate that the downwards bias is much too high. In the sample by Malul (2021) with a true wage gap g close to $g = 0.3$, men underestimate the gender wage gap by about 36% and women only by about 22%. This is, of course, far from the at least about 50% downwards bias that is, given our simulation results, necessary for underprivileged agents to be on average closer to reality than the privileged, in line with stylised fact (ii) and also consistent with our back-of-the-envelope calibration for overall inequality perceptions. This might indicate that, in contrast to perceptions of *overall* inequality, perceptions of *gender* inequality might be formed as a combination of global and local signals. We calibrate our model in line with the composite rule in Eq. (4) in subsection "Individual perception" to spell out the implications of this in more detail.

Disaggregating local perceptions

The relatively straightforward impact of income and homophily on the perceived wage gaps at the aggregate level masks substantial differences for individual agents. Here, *neighbourhood disassortativity D* presents the critical link between an agent's income and wage gap perception. It is the rate of link neighbours of the opposite privilege class. Hence, $D_i = 0$ means that agent i is located within a segregated community of one's own privilege class whereas i being isolated among members of the other class results in $D_i = 1$; a disassortativity slightly below 0.5 (the exact rate depending on the size of the ego network with 0.33 for 5 link neighbours) indicates equal numbers within the group.

We expect disassortativity to be positively correlated to income for underprivileged agents: The wage gap decreases diversity among the poorest of them because they become isolated from the poorest privileged agents; concurrently, those underprivileged agents with a pre wage gap income higher than the global median now get pulled closer to it and experience a high diversity level (with disassortativity close to 0.5) due to the low selectivity there. Contrary to that, we expect a negative relationship between income and disassortativity when it comes privileged agents: The highest earners of them are also the richest overall and they experience few underprivileged agents in their link-neighbourhoods amounting to rather low disassortativity levels. By contrast, the wage

---

[8] For this calibration, we calculate the mean perceived Gini coefficient for the whole sample population and vary ρ to replicate the maximum and minimum perceived Gini that Choi (2019) finds for his large sample from the International Social Survey Program. The sample covers 32 OECD countries and 30 sample years. Over all countries, he finds a minimum perceived Gini of $G_{min} = 0.1276$, a mean perceived Gini of $G_{mean} = 0.1708$ and a maximum perceived Gini of $G_{max} = 0.2534$. The maximum Gini G_{max} translates to $\rho \approx 4$, while the minimum Gini G_{min} implies $\rho \approx 14$.

Fig. 4 Heatmap of wage gap perceptions (as a ratio to the correct wage gap *g*) for varying privilege, disassortativity *D* (rate of links to the other privilege class) and income *Y* for $\rho = 8$ and varying actual wage gaps $g \in \{0.1; 0.2; 0.3; 0.4\}$. Upper panels for the underprivileged agents, lower panels for the privileged ones

Fig. 5 Heatmap of wage gap perceptions (as a ratio to the correct wage gap *g*) for varying privilege, disassortativity *D* (rate of links to the other privilege class) and income *Y* for $g = 0.2$ and varying actual wage gaps $\rho \in \{1; 4; 8; 14\}$. Upper panels for the underprivileged agents, lower panels for the privileged ones

gap pushes the poorest privileged closer to the the overall median income meaning more frequent encounters with underprivileged agents as explained above. One would expect these mechanisms to grow more powerful in *g*, for this results in a greater divergence of in-group income ranks and global ones. The homophily strength $\rho > 0$ should work as another catalyst of the described mechanisms.

Figures 4 and 5 depict this relationship for examples of *g* and ρ.[9] A vertical section through a heatmap represents all levels of neighbourhood disassortativity—and resulting perceptions—that any specific individual agent (identified across MC runs by their income) for the respective combination of wage gap and homophily. Chance has a substantial impact on low income ranks, and the disassortativity varies largely between MC runs. Due to the linkage selectivity, the range of realised disassortativity levels shrinks for higher income ranks. That this effect becomes stronger with rising ρ testifies to the

---

[9] The figures present a selection of values for *g* and ρ that are empirically plausible. Yet, the data for other parameter combinations display the complex and non-monotonous patterns as well.

linkage procedure discussed in the section on the model setup. Thus, underprivileged agents with higher incomes only display high disassortativity levels, in line with the theoretical expectations above. However, while the range of realised disassortativity levels shrinks in income for the privileged, where this range lies for the highest incomes depends on *g*. Furthermore, the relationship between income and disassortativity is highly non-linear.

Non-linear and even non-monotonous patterns also surface in the effects that a rising ρ for a given *g*—or vice versa—have on disassortativity. An agent's overall income rank changes in *g*, and the underprivileged agents with the highest income can grow less disassortative for very high *g* as Fig. 4 illustrates. Figure 5 highlights that even moderate homophily levels significantly impact disassortativity, especially of privileged agents. Nevertheless, while an increase in $\rho > 0$ from low to moderate levels decreases the range of realised disassortativity levels, this range may expand again for some agents when ρ gets high.

An agent's different neighbourhood disassortativity levels themselves affect this agent's local wage gap perception (*l*). With this, Figs. 4 and 5 testify to the interaction of *Y* and *D* (both values being themselves traced to the pre wage gap distribution, *g*, and ρ) and its complex impact on *l*. Namely, they point to three obvious phenomena:

- For a given ρ, a higher actual wage gap *g* does not imply a higher perceived wage gap for all agents: Trivially, an agent's perception depends on their link-neighbours but a higher actual wage gap for the whole population can mean shrinking local differences between privileged and underprivileged members of some ego networks.
- Given any income and income rank distribution implied by the wage gap, rising homophily lowers the average wage gap perceptions overall as expected. However, some agents' *l* behaves non-monotonically and may rise in ρ.
- Chance has a significant impact on many agents' ego networks (i.e., their disassortativity) and the resulting perceptions, as the great variety in *D* and *l* between MC runs—for vertical sections through heatmaps—demonstrates.

As a corollary of these phenomena, wage gap perception differs from general inequality one in Schulz et al. (2022), one cannot predict an agent's individual wage gap estimate *l* by their income rank; and their privilege class does not work as a predictor, too. It is neither *Y* nor *D* alone that determines an agents' wage gap perception but the interaction of the two. Given sufficiently high levels of *g* and ρ, paints a clear and monotonous picture for high income ranks and at the edges of the realised disassortativity levels. However, for the majority of agents and MC runs (i.e. disassortativity levels of an agent with a given income rank), there is no such monotonous behaviour of *l*.

These behaviours occur because the likeliness of an agent with a given income rank in their privilege class to be picked as another agent's link-neighbour simultaneously depends on *g* and ρ, as shown in section "Model". The wage gap determines the likeliness ranking of agents as link neighbours. If for one agent the relative likeliness ranks of only two other agents, who belong to different privilege classes, change, all three agents may experience notably changed wage gap perceptions. The homophily level specifies how much the likelihood of selection as link neighbours differs between agents. Hence,

Fig. 6 Implied weights for the global signal for both privilege classes and the empirical calibration
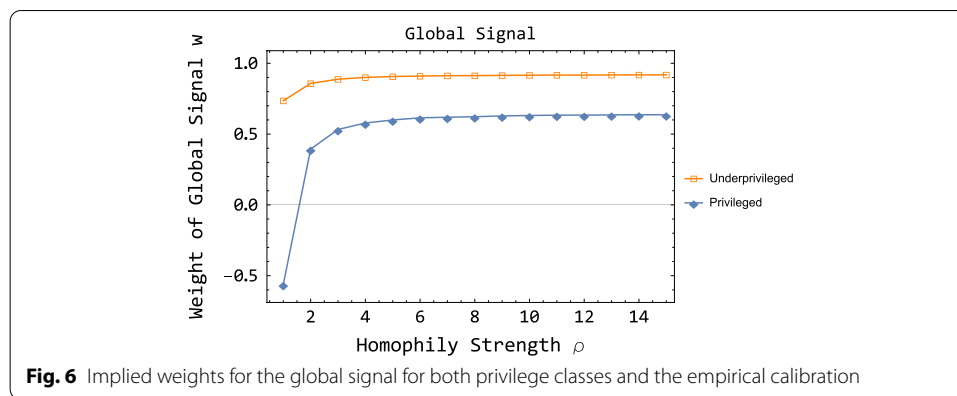
the perception landscapes become more 'rugged' when ρ increases. Moreover, like disassortativity, it depends on income how much the perceptions vary between MC runs. Thereby, agents with the highest incomes exhibit the least variability because they are the most selective regarding their link-neighbours, giving less room for chance.

Empirical calibration and composite signal

The simulation results in subsection "Mean local perceptions" indicate that the effect strength from homophilic segregation on local perceptions is too strong to replicate the empirically observed effect sizes in Malul (2021). To account for this, we extend our rule for information processing along the lines introduced in subsection "Individual perception" to allow for a *composite signal* aggregated from global and local information, respectively. We calibrate our baseline model with the empirically observed parameters to deduce the weight both the privileged (male) and the underprivileged (female) class puts on the global signal. To achieve this, we use a 'naive estimator', where we replace the individual p_i and l_i in Eq. (4) with the respective sample means $\bar{p}$ and $\bar{l}$. This estimator is exact, whenever the individual weights w_i that individuals i place on the global signal are uncorrelated to local perceptions l_i, as we show in "Appendix". The empirical gender wage gap for the Israeli sample is $g = 0.317$, with a female labour force participation rate of 0.472 and a male one of approximately 0.528. The target mean perception of the underprivileged is $\bar{p}_U = 0.292$ and of the privileged $\bar{p}_P = 0.203$. We initialise our model with a total population of $N = 1000$ which implies an underprivileged population size of $N_U = 472$ and a privileged population size of $N_P = 528$. The estimated implied mean weights $\bar{w}_U$ and $\bar{w}_P$ for both the underprivileged and privileged class and different ρ are shown in Fig. 6.

For the whole range of estimates, we find that the underprivileged group places a much higher weight on the global signal than the privileged one. In the (empirically implausible) case of $\rho = 1$, the implied weight of the privileged even becomes negative, which would amount to the privileged part of the population actively discarding global information in favour of 'overweighting' their biased local signals. The gap between the underprivileged and privileged classes becomes more narrow with higher homophily when the local perceptions of the underprivileged decrease but continues to be sizeable for all considered ρ. Note that this also implies that the composite perceptions of the underprivileged are much less volatile, as they put more weight on the (correct) global
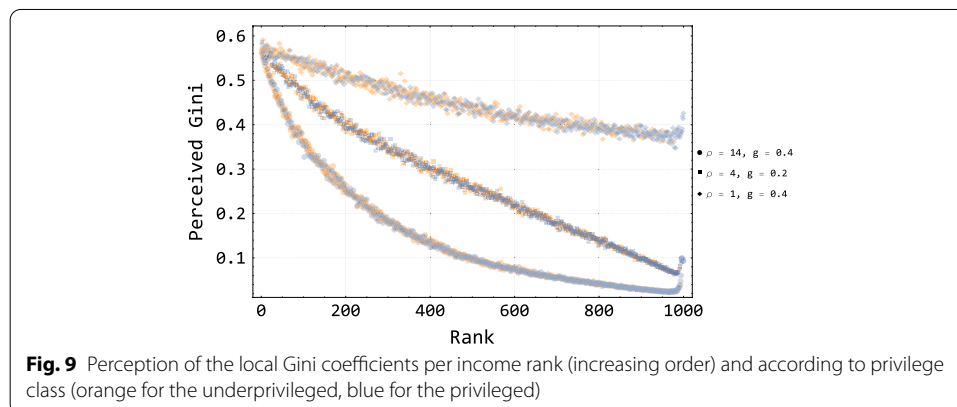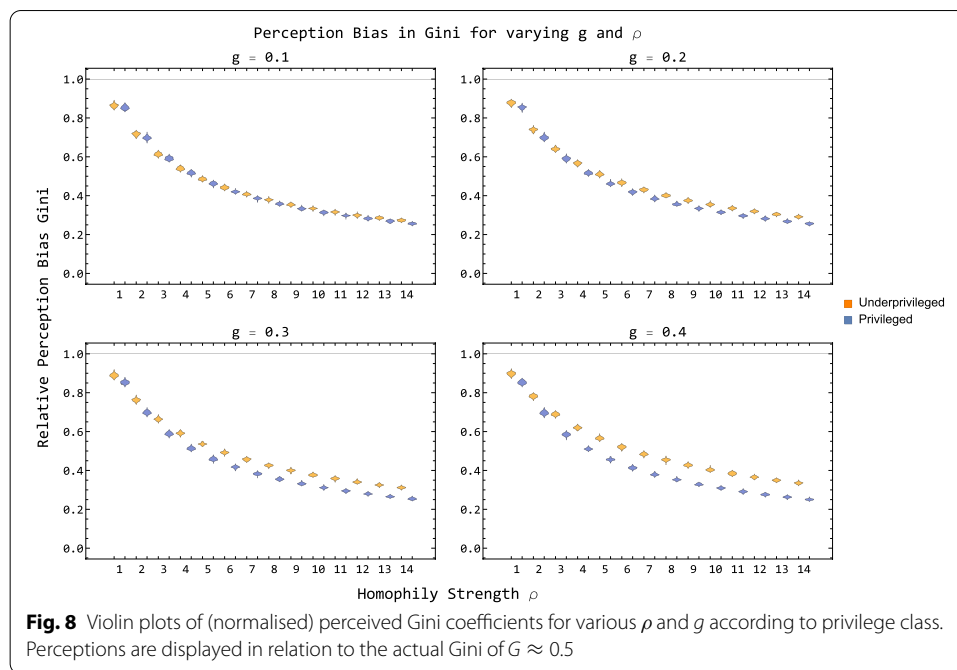
Fig. 7 Violin plots for the composite perceptions of privileged and underprivileged agents for the empirical calibration and $\rho = 8$
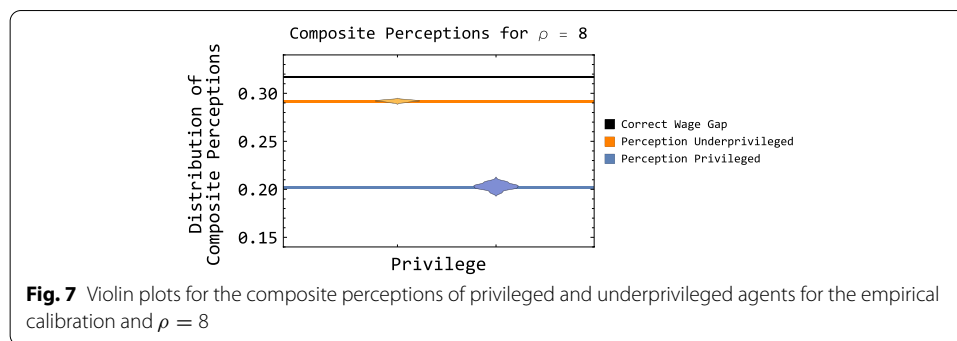
signal rather than the (noisy) local one, which is evident from the exemplary violin plots for $\rho = 8$ in Fig. 7. In this case of $\rho = 8$, the difference in the variance of estimates is more than one order of magnitude, with the variance of mean composite perceptions of the privileged being more than tenfold the variance of mean composite perceptions of the underprivileged.

Our calibration exercise thus yields two main results: Firstly, we find that local perception formation on skewed information sets is insufficient to generate the observed empirical effect sizes. Only for composite signals that combine both a (correct) global signal with the local perception can the model be reconciled with the empirical evidence provided in Malul (2021). By contrast, the local signal is sufficient for perceptions of *global inequality* (Schulz et al. 2022), hinting at conceptual differences between both. They are important because they might imply that the concept of a 'wage gap' can be much more easily conveyed within educational campaigns than the arguably much more complex and ambiguous concept of 'inequality'. Indeed, the empirical evidence suggests that concept complexity and the associated cognitive costs might exhibit significant effects on individual behaviour (Oprea 2020). Secondly, we find that the underprivileged place much higher weight on the global signal than the privileged, with much less noisy estimates. This is consistent with the empirical finding that the adversely affected part of the population is more interested in global information about the issue, as is also documented on the micro level within the empirical psychological literature (Wu 2021).

Our results from calibrating the composite perception rule suggest that belief formation regarding intergroup inequality is based not only on local but also on global information. Since our proposed mechanism is purely based on economic considerations, this also indicates that the implied differential weights men and women place on the global signal requires an explanation by factors outside of the network-segregation mechanism we propose.

The impact of privilege on general inequality perceptions

In contrast to perceptions of the wage gap, the ones of biases regarding overall rather than gender-based inequality can be generated from network segregation only, as Fig. 8 demonstrates. For all considered $\rho > 0$ and $g > 0$, the underprivileged appear to observe a greater degree of economic inequality within the population. The perceived degree of inequality obviously decays in ρ as network segregation increases. Hence, the introduction of privilege classes to the model does not alter its outcomes concerning general

Fig. 8 Violin plots of (normalised) perceived Gini coefficients for various ρ and g according to privilege class. Perceptions are displayed in relation to the actual Gini of $G \approx 0.5$

Fig. 9 Perception of the local Gini coefficients per income rank (increasing order) and according to privilege class (orange for the underprivileged, blue for the privileged)

inequality perceptions described by Schulz et al. (2022). The relative bias of the privileged increases in *g*, that is, the underprivileged become relatively more accurate in their perceptions compared to the privileged, the higher the wage gap is. Paradoxically, being economically underprivileged thus translates into privileged epistemic access regarding the actual state of inequality in our model.

In contrast to the vast literature on redistributive preferences according to race and gender surveyed in section "Related literature", this finding constitutes an emergent outcome of the homophily in ego network formation. As Schulz et al. (2022) note, the homophilic graph formation lets the income differences in the numerator of Eq. (5) increase linearly at most in income rank. In contrast, the mean perceived income in the denominator increases much faster exponentially due to the underlying exponential wage distribution. As a result, individuals with lower income (ranks) perceive higher Gini coefficients *L* and are more accurate on average. Pushing the underprivileged to the

lower income ranks by the wage gap g thus increases their average perception. Figure 9 highlights this effect, with perceived Ginis L visibly decaying in income rank. As is readily apparent from the same figure, the density of underprivileged agents (with orange markers) is much higher at lower ranks of the rank distribution, indicating that those also perceive higher levels of inequality on average.

We are thus able to replicate the finding that perceptions anchored within ego networks (L) are sufficient to generate the empirically observed patterns of overall inequality perceptions. Our findings also directly suggest that it is not necessarily non-economic factors like normative beliefs or self-interest driving the gender and racial differentials in support of redistributive policies. Instead, it might simply be the case that skewed samples from homophilic ego networks render the privileged particularly ignorant regarding issues of economic inequality, as our simulation results would suggest.

Discussion

The agent-based simulation model presented in this paper suggests that the mild underestimation of gender and racial wage gaps, which empirical studies describe, results from individuals processing information obtained by combining one's link-neighbours' incomes with a global signal reflecting the actual wage gap. However, underprivileged agents put considerably higher weights on the global signal than privileged ones. Therefore, those adversely affected by a wage gap may perceive it more severely, i.e., closer to its actual severity, simply because they listen more carefully to global information. Likewise, underprivileged agents underestimate income inequality, in general, less strongly than privileged ones because of their lower selectivity when choosing link neighbours, the latter itself being a corollary of an agent's own income level and rank. Different from wage gap perception, a global signal is not necessary to explain the empirical underestimation of general wage inequality. Furthermore, the level of underestimation hereby depends solely on an individual's income (with lower incomes typically meaning a less severe underestimation)—not their privilege class.

We achieve these conclusions by imposing a rather strong identification assumption, i.e., local perceptions and weights for the two signals being uncorrelated or homogeneous within groups. This leads to a 'naive estimator' for the implied weights for local and global signals, where we simply substitute the mean local perceptions $\bar{l}_U$ and $\bar{l}_P$ for the underprivileged and privileged in Eq. (4) that is stated in terms of the individual local perceptions l_i. Fortunately, we recover the mean weights of the respective subpopulations $\bar{w}_U$ and w_P, whenever the individual cross-sectional weights and local perceptions are uncorrelated for the given subpopulation or whenever weights are homogeneous in them, as we show in "Appendix". For a positive correlation between individual weights of the global signal and local perceptions, we underestimate the true mean and for a negative correlation, we overestimate the mean weight with our naive estimator. One could make a plausible case for both types of correlation: Either high local perceptions let individuals pay more attention to their immediate lived experience, which would imply a negative correlation or local perceptions of high inequality prime people into being more responsive to global signals about the issue which would then imply a positive relationship between both. Since we do not have any information on the correlation structure within the sample by Malul (2021), this assumption of zero correlation appears to be the

most natural one to us, then directly implying the naive estimator. For further applied studies, working out the actual correlational patterns should, however, be a priority and might shed light on further theoretical implications of the two correlation mechanisms we sketched, i.e., correlation patterns differing between privileged and underprivileged groups, therefore, explaining part of the difference in estimated implied weights.

For policy-makers and the public debate, this finding suggests that underprivileged groups are a reliable source, and their account of the wage gap is not strongly biased by subjective experience or even self-interest but on a higher emphasis on what is objectively the case. Furthering such emphasis on the globally true value, e.g. through education treatments, seems to be especially promising in the context of a wage gap that is easy to comprehend. These measures should especially target the privileged who put lower weight on the global signal and rely more on the biased conclusions that they draw from observing their ego networks. Another challenge raised by our model is that income homophily can perpetuate homophily in privilege class (e.g., racist self-segregation) even though individuals do not willingly segregate by these privilege groups. Thus, when taking any actions to combat such segregation, policy makers should observe income homophily, too.[10]

In many ways, the global signal and the associated weights remain a residual quantity within our model that is not further explained. In obvious analogy to the concept of "total factor productivity" as a growth residual in economics, the global weights we identify are thus more than anything "measure[s] of our ignorance" Abramovitz (1993) concerning the actual underlying cognitive processes. However, we hope that our taxonomy of global and local signals as well as the mechanism for composite perception formation informs further empirical research on determinants of the global weight. In particular, we identified the correlational structure between local perceptions and global weights to be of potentially immense importance in determining composite perceptions.

We find that individual level phenomena react highly sensitively to both predetermined and random initial conditions. This property potentially poses obstacles to applied empirical research—given resemblance between our model network and the real world one. The complexity of the landscape indicates that average population findings mask important heterogeneity on an individual level; thus there is no simple meaningful aggregate representation. Moreover, our findings indicate that the relationships between income/actual wage gap, disassortativity, and perceptions do not only depend on the actual wage gap, the strength of homophily and the privilege class. Instead, the ego network disassortativity and perception vary along the income distribution—and the steepness of this variability is also non-monotonous. Therefore, it is impossible to fit any parsimonious monotonic function to any of these relationships.[11] Per construction, the relationships are deterministic in our model, but in the presence of the sketched non-monotonicities, empirical work might fail to detect the underlying structures, relationships and their causal mechanisms we impose. The striking non-monotonicities in individual perceptions imply a need for great caution when employing

---

[10] We thank an anonymous reviewer for raising this last policy implication.

[11] Indeed, the Pearson correlation coefficients between local perceptions l and disassortativity D for all simulation runs and considered parameter combinations do not exceed 0.1 in absolute value.

standard statistical tools in such further research, though. Namely, Pearson correlations are not meaningful since the observed relationships are highly case-dependent. Moreover, standard measures of non-linear relationships like Kendall's τ or Spearman's ρ do not work in the studied situation of two groups that react differently to a change in an external variable. Finally, individual perceptions depend on the specific network layout. In turn, this layout emerges from the interplay of the actual wage gap, homophily, and chance. Hence, reliable predictions are impossible at an aggregate level and instead require investigation of individuals in their network interactions.

Our general framework points to obvious theoretical extensions: First, we assumed a common ρ for all individuals throughout all experiments, irrespective of their privilege class as the most parsimonious assumption. A natural extension would allow for heterogeneity in ρ across groups to examine differential homophily. Also, we only considered two groups each so far but neglected that identity is an inherently multidimensional construct with discrimination and marginalisation operating on different layers. Therefore, model extensions should consider the interaction of both racial and gender-based discrimination on individual perceptions. An *intersectional* perspective should account for the fact that wage gaps are *superadditive* (Stelzner and Bahn 2020; Bright et al. 2016). Apart from considering more than two groups, exploring differing functional forms for the superadditivity of wage gaps will surely also provide interesting and promising new avenues for further research.

Appendix

Calibrating the composite wage gap

Since we only observe the population means in our sample and especially are not able to determine any cross-sectional correlation between our variables, we impose some conditions on our estimator. In particular, we use the following expression as our naive estimator $\hat{w}$ for the implied weight a subpopulation (the privileged or underprivileged) places on the global signal, where we replace the individual p_i and l_i in Eq. (4) with the respective sample means $\bar{p}$ and $\bar{l}$, i.e.,

$$\hat{w} = \frac{\bar{p} - \bar{l}}{g - \bar{l}}, \tag{6}$$

where $\bar{p}$ corresponds to the mean perceived wage gap and $\bar{l}$ the mean local signal of the given sampled subpopulation as well as g as the true global wage gap. Assume that all individuals form perceptions according to our composite perception rule. The estimator $\hat{w}$ then corresponds to the true mean weight of the sampled subpopulation, whenever either i) w is equal for all considered individuals within the subpopulation or ii) w and l and thus $(1 - w)$ and l are uncorrelated. To see this, consider the composite perception rule for a given individual i:

$$p_i = w_i \cdot g + (1 - w_i) \cdot l_i. \tag{7}$$

Averaging over all individuals i within the considered subpopulation of size n yields

$$n^{-1} \cdot \Sigma_{i=1}^{n} p_i = n^{-1} \cdot \Sigma_{i=1}^{n} w_i \cdot g + n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i)l_i. \tag{8}$$

Expressing this in the bar notation from above in Eq. (6), we get

$$\bar{p} = \bar{w}g + n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i)l_i. \tag{9}$$

From expression (7), it becomes obvious that there is one degree of freedom left in our specification: The relationship between the vector of $(1 - w)_{i=1}^{n}$ and $(l)_{i=1}^{n}$. We consider two cases:

(i) Assume first that $w_i = \bar{w}, \forall i = 1, 2, ..., n$, that is, all individuals place equal weights on the global signal. This also implies that $n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i) = (1 - \bar{w})$ and thus allows to rewrite (7) as

$$\hat{w} = \frac{\bar{p} - \bar{l}}{g - \bar{l}}, \tag{10}$$

$$\bar{p} = \bar{w}g + (1 - \bar{w})\bar{l} \tag{11}$$

which is the desired result and gives rise to $\bar{w} = \hat{w}$ according to Eq. (6).

(ii) Assume now that the sample correlation between l and $(1 - w)$, $r_{1-w,l}$ exists and is exactly zero. Recall the definition of a sample correlation as

$$r_{wl} = \frac{\sum_{i=1}^{n}(1 - w_i)l_i - n \cdot \overline{(1 - w)} \cdot \bar{l}}{n \, s'_{1-w} \, s'_l} = \frac{\sum_{i=1}^{n}(1 - w_i)l_i - n \cdot (1 - \bar{w}) \cdot \bar{l}}{n \, s'_{1-w} \, s'_l}, \tag{12}$$

due to $\overline{(1 - w)} = (1 - \bar{w})$ as well as $s'_l > 0$ and $s'_{1-w} > 0$ as the (corrected) sample standard deviation, i.e., both vectors exhibiting some variability. Imposing $r_{1-w,l} = 0$ implies immediately that

$$\sum_{i=1}^{n}(1 - w_i)l_i = n \cdot (1 - \bar{w}) \cdot \bar{l} \quad \text{and thus}$$

$$n^{-1} \sum_{i=1}^{n}(1 - w_i)l_i = (1 - \bar{w}) \cdot \bar{l}. \tag{13}$$

This again allows us to rewrite Eq. (7) as

$$\bar{p} = \bar{w}g + (1 - \bar{w})\bar{l} \tag{14}$$

which is the desired result and gives rise to $\bar{w} = \hat{w}$ according to Eq. (6).

The estimator is thus only exact, whenever either the correlation between $(1 - w)_{i=1}^{n}$ and $(l)_{i=1}^{n}$ is exactly zero or if all individuals place exactly the same weights on the global and local signals, respectively. Whenever $r_{1-w,l} > 0$, $\hat{w} > \bar{w}$ and we thus overestimate the true population mean (and vice versa). This follows straightforwardly from Eq. (7). Solving for $\bar{w}$, the equation implies that, generally,

$$\bar{w} = \frac{\bar{p} - n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i)l_i}{g}. \tag{15}$$

The estimator implies

$$\hat{w} = \frac{\bar{p} - (1 - \bar{w})\bar{l}}{g}. \tag{16}$$

Whenever $r_{1-w,l} > 0$ it follows that $n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i) > (1 - \bar{w})\bar{l}$ by Eq. (12) which implies by Eqs. (15) and (16) that $\bar{w} < \hat{w}$. Conversely, $r_{1-w,l} < 0$ implies by the same token that the inequality reverses and $\bar{w} > \hat{w}$, since $n^{-1} \cdot \Sigma_{i=1}^{n}(1 - w_i) < (1 - \bar{w})\bar{l}$. Note finally that since $r_{1-w,l} < 0$ implies that $r_{w,l} > 0$ must hold, $r_{w,l} > 0$ implies that $\bar{w} > \hat{w}$ (and vice versa).

Declarations

Ethics approval and consent to participate
The study reported in the manuscript involved no human participants, human data or human tissue.

Consent for publication
The manuscript does not contain any individual person's data.

Competing interests
The authors declare that they have no competing interests.

Author details
[1]Institute for Political Science, University of Bamberg, Bamberg, Germany. [2]Economics Department, University of Bamberg, Bamberg, Germany.

References
Abramovitz M (1993) The search for the sources of growth: areas of ignorance, old and new. J Econ Hist 53(2):217–243
Akchurin M, Lee C-S (2013) Pathways to empowerment: repertoires of women's activism and gender earnings equality. Am Sociol Rev 78(4):679–701
Akerlof GA, Yellen JL (1990) The fair wage-effort hypothesis and unemployment. Q J Econ 105(2):255–283
Alesina A, Giuliano P (2011) Preferences for redistribution. In: Benhabib J, Bisin A, Jackson MO (eds) Handbook of social economics, vol 1. Elsevier, Amsterdam, pp 93–131
Alesina A, Ferroni MF, Stantcheva S (2021) Perceptions of racial gaps, their causes, and ways to reduce them. NBER Working Paper Series 29245
Andreoni J, Vesterlund L (2001) Which is the fair sex? Gender differences in altruism. Q J Econ 116(1):293–312
Auspurg K, Hinz T, Sauer C (2017) Why should women get less? Evidence on the gender pay gap from multifactorial survey experiments. Am Sociol Rev 82(1):179–210
Banerjee A, Yakovenko VM (2010) Universal patterns of inequality. New J Phys 12(7):075032
Banerjee A, Yakovenko VM, Di Matteo T (2006) A study of the personal income distribution in Australia. Physica A 370(1):54–59
Barde S, Van Der Hoog S (2017) An empirical validation protocol for large-scale agent-based models. Bielefeld Working Papers in Economics and Management

Bavetta S, Li Donni P, Marino M (2019) An empirical analysis of the determinants of perceived inequality. Rev Income Wealth 65(2):264–292

Boguná M, Pastor-Satorras R, Díaz-Guilera A, Arenas A (2004) Models of social networks based on social distance attachment. Phys Rev E 70(5):06122

Bright LK, Malinsky D, Thompson M (2016) Causally interpreting intersectionality theory. Philos Sci 83(1):60–81

Carter ER, Brady ST, Murdock-Perriera LA, Gilbertson MK, Ablorh T, Murphy MC (2019) The racial composition of students' friendship networks predicts perceptions of injustice and involvement in collective action. J Theor Soc Psychol 3(1):49–61

Ceccardi T (2021) How do beliefs about the gender wage gap affect the demand for public policy? Differences 2021:14

Chandra A (2000) Labor-market dropouts and the racial wage gap: 1940–1990. Am Econ Rev 90(2):333–338

Choi G (2019) Revisiting the redistribution hypothesis with perceived inequality and redistributive preferences. Eur J Polit Econ 58:220–244

Cusack T, Iversen T, Rehm P (2006) Risks at work: the demand and supply sides of government redistribution. Oxf Rev Econ Policy 22(3):365–389

Dall J, Christensen M (2002) Random geometric graphs. Phys Rev E 66(1):016121

Davidai S, Walker J (2021) Americans misperceive racial disparities in economic mobility. Personal Soc Psychol Bull 48(5):793–806

Decker R, Haltiwanger J, Jarmin R, Miranda J (2014) The role of entrepreneurship in us job creation and economic dynamism. J Econ Perspect 28(3):3–24

Derzsy N, Néda Z, Santos M (2012) Income distribution patterns from a complete social security database. Physica A 391(22):5611–5619

Dos Santos PL (2017) The principle of social scaling. Complexity: 8358909

Dragulescu A, Yakovenko VM (2001) Evidence for the exponential distribution of income in the USA. Eur Phys J B Condens Matter Complex Syst 20(4):585–589

Dragulescu A, Yakovenko VM (2001) Exponential and power-law probability distributions of wealth and income in the United Kingdom and the United States. Physica A 299(1–2):213–221

Fagiolo G, Guerini M, Lamperti F, Moneta A, Roventini A (2019) Validation of agent-based models in economics and finance. In: Beisbart C, Saam NJ (eds) Computer simulation validation: fundamental concepts, methodological frameworks, and philosophical perspectives. Springer, Cham, pp 763–787

Falk A, Knell M (2004) Choosing the joneses: endogenous goals and reference standards. Scand J Econ 106(3):417–435

Fixman M, Berenstein A, Brea J, Minnoni M, Travizano M, Sarraute C (2016) A Bayesian approach to income inference in a communication network. In: 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM). IEEE, pp 579–582

Foley S, Kidder DL, Powell GN (2002) The perceived glass ceiling and justice perceptions: an investigation of Hispanic law associates. J Manag 28(4):471–496

Franke R, Westerhoff F (2017) Taking stock: a rigorous modelling of animal spirits in macroeconomics. J Econ Surv 31(5):1152–1182

Furnham A, Wilson E (2011) Gender differences in estimated salaries: a UK study. J Socio-Econ 40(5):623–630

García-González J, Forcén P, Jimenez-Sanchez M (2019) Men and women differ in their perception of gender bias in research institutions. PLoS ONE 14(12):0225763

Gräbner C (2018) How to relate models to reality? An epistemological framework for the validation and verification of computational models. J Artif Soc Soc Simul 21(3):8

Grüne-Yanoff T (2009) Learning from minimal economic models. Erkenntnis 70(1):81–99

Grüne-Yanoff T, Verreault-Julien P (2021) How-possibly explanations in economics: anything goes? J Econ Methodol 28(1):114–123

Haaland I, Roth C (2021) Beliefs about racial discrimination and support for pro-black policies. Rev Econ Stat (forthcoming)

Hamilton MJ, Milne BT, Walker RS, Burger O, Brown JH (2007) The complex structure of hunter-gatherer social networks. Proc R Soc B Biol Sci 274(1622):2195–2203

Hampton MB, Heywood JS (1993) Do workers accurately perceive gender wage discrimination? ILR Rev 47(1):36–49

Harting P, Radi D (2020) Residential segregation: the role of inequality and housing subsidies. J Econ Behav Organ 178:801–819

Hauser OP, Norton MI (2017) (Mis)perceptions of inequality. Curr Opin Psychol 18:21–25

Hu X, Liang C-Y (2022) Does income redistribution prevent residential segregation? J Econ Behav Organ 193:519–542

Jagielski M, Kutner R (2013) Modelling of income distribution in the European Union with the Fokker-Planck equation. Physica A 392(9):2130–2138

Karlsson N, Gärling T, Dellgran P, Klingander B (2005) Social comparison and consumer behavior: when feeling richer or poorer than others is more important than being so. J Appl Soc Psychol 35(6):1206–1222

Kets W, Sandroni A (2019) A belief-based theory of homophily. Games Econ Behav 115:410–435

Kinder DR, Winter N. Exploring the racial divide: blacks, whites, and opinion on national policy. Am J Polit Sci. 2001;45(2):439–56.

Kraus MW, Rucker JM, Richeson JA (2017) Americans misperceive racial economic equality. Proc Natl Acad Sci 114(39):10324–10331

Kraus MW, Onyeador IN, Daumeyer NM, Rucker JM, Richeson JA (2019) The misperception of racial economic inequality. Perspect Psychol Sci 14(6):899–921

Krieger N (1992) Overcoming the absence of socioeconomic data in medical records: validation and application of a census-based methodology. Am J Public Health 82(5):703–710

Kunze A (2018) The gender wage gap in developed countries. In: Averett SL, Argys LM, Hoffman SD (eds) The Oxford handbook of women and the economy. Oxford University Press, Amsterdam, pp 369–394

Larson JM (2017) The weakness of weak ties for novel information diffusion. Appl Netw Sci. https://doi.org/10.1007/s41109-017-0034-3

Leo Y, Fleury E, Alvarez-Hamelin JI, Sarraute C, Karsai M (2016) Socioeconomic correlations and stratification in social-communication networks. J R Soc Interface 13(125):20160598

Linos K, West M (2003) Self-interest, social beliefs, and attitudes to redistribution. Re-addressing the issue of cross-national variation. Eur Sociol Rev 19(4):393–409

Mac Carron P, Kaski K, Dunbar R (2016) Calling Dunbar's numbers. Soc Netw 47:151–155

Mäki U (2009) MISSing the world. Models as isolations and credible surrogate systems. Erkenntnis 70:29–43

Malul M (2021) (Mis)perceptions about the gender gap in the labor market. In: Forum for social economics. Taylor & Francis, pp 1–9

Manski C (1981) Structural analysis of discrete data with econometric applications. MIT Press, Cambridge

Mayerhoffer DM, Schulz J (2022) Marginalisation and misperception: perceiving gender and racial wage gaps in ego networks. In: Benito RM, Cherifi C, Cherifi H, Moro E, Rocha LM, Sales-Pardo M (eds) Complex networks & their applications X. Springer, Cham, pp 768–779

McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: homophily in social networks. Ann Rev Sociol 27(1):415–444

Milli L (2021) Opinion dynamic modeling of news perception. Appl Netw Sci. https://doi.org/10.1007/s41109-021-00412-4

Morgan J, Kelly NJ (2017) Social patterns of inequality, partisan competition, and Latin American support for redistribution. J Polit 79(1):193–209

Newman ME (2009) Random graphs with clustering. Phys Rev Lett 103(5):058701

Newman ME, Strogatz SH, Watts DJ (2001) Random graphs with arbitrary degree distributions and their applications. Phys Rev E 64(2):026118

Nirei M, Souma W (2007) A two factor model of income distribution dynamics. Rev Income Wealth 53(3):440–459

Oancea B, Andrei T, Pirjol D (2017) Income inequality in Romania: the exponential-Pareto distribution. Physica A 469:486–498

Oprea R (2020) What makes a rule complex? Am Econ Rev 110(12):3913–3951

Pfeifer C, Stephan G (2019) Why women do not ask: gender differences in fairness perceptions of own wages and subsequent wage growth. Camb J Econ 43(2):295–310

Quadagno J, Blekesaune M (2003) Public attitudes towards welfare state policies. Eur Sociol Rev 19(5):415–427

Schneider M (2015) Revisiting the thermal and superthermal two-class distribution of incomes. Eur Phys J B 88(1):1–10

Schulz J, Mayerhoffer D, Gebhard A (2022) A network-based explanation of inequality perceptions. Soc Netw 70:306–324

Shaikh A (2017) Income distribution, econophysics and Piketty. Rev Polit Econ 29(1):18–29

Shaikh A, Jacobo JE (2020) Economic arbitrage and the econophysics of income inequality. Rev Behav Econ 7:299–315

Shaikh A, Papanikolaou N, Wiener N (2014) Race, gender and the econophysics of income distribution in the USA. Physica A 415:54–60

Silva AC, Yakovenko VM (2004) Temporal evolution of the thermal and super thermal income classes in the USA during 1983–2001. Europhys Lett 69(2):304

Smith JA, McPherson M, Smith-Lovin L (2014) Social distance in the United States. Am Sociol Rev 79(3):432–456

Stelzner M, Bahn K (2020) Discrimination and monopsony power. Rev Black Polit Econ 49(2):152–174

Talaga S, Nowak A (2019) Homophily as a process generating social networks: insights from social distance attachment model. J Artif Soc Soc Simul 23(2):6

Tao Y, Wu X, Zhou T, Yan W, Huang Y, Yu H, Mondal B, Yakovenko VM (2019) Exponential structure of income inequality: evidence from 67 countries. J Econ Interact Coord 14(2):345–376

Törnqvist L, Vartia P, Vartia YO (1985) How should relative changes be measured? Am Stat 39(1):43–46

Virtanen M, Kivimäki M, Kouvonen A, Elovainio M, Linna A, Oksanen T, Vahtera J (2007) Average household income, crime, and smoking behaviour in a local area: the Finnish 10-town study. Soc Sci Med 64(9):1904–1913

Watts DJ, Strogatz SH (1998) Collective dynamics of small-world networks. Nature 393(6684):440–442

Williams MJ, Elizabeth LP, Spencer-Rodgers J (2010) The masculinity of money: automatic stereotypes predict gender differences in estimated salaries. Psychol Women Q 34(1):7–20

Wu K (2021) Invisibility of social privilege to those who have it. In: Academy of Management Proceedings 2021(1), 10776 Academy of Management Briarcliff Manor, NY, 10510

Xu Y, Santi P, Ratti C (2021) Beyond distance decay: discover homophily in spatially embedded social networks. Ann Am Assoc Geogr 112(2):505–521

Zhou W-X, Sornette D, Hill RA, Dunbar RI (2005) Discrete hierarchical organization of social group sizes. Proc R Soc B Biol Sci 272(1561):439–444

Publisher's Note