Schulz-Gebhard, Jan; Mayerhoffer, Daniel M

A Network Approach to Expenditure Cascades

Date of secondary publication: 07.04.2024

Accepted Manuscript (Postprint), Article

Persistent identifier: urn:nbn:de:bvb:473-irb-923734

Primary publication

Schulz-Gebhard, Jan; Mayerhoffer, Daniel M. (2023): A Network Approach to Expenditure Cascades, in: Review of Behavioral Economics. Hanover, Mass. : Now Publishers, DOI: 10.1561/105.00000174.

A Network Approach to Expenditure Cascades*

Accepted Manuscript

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

[1]Department of Economics, University of Bamberg

[2]Faculty of Social and Behavioural Sciences, University of Amsterdam

Abstract

*The final publication is available from now publishers via http://dx.doi.org/10.1561/105.00000174. Daniel Mayerhoffer's contribution was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 430621735. Furthermore, financial support for the paper by the University of Bamberg through the Fres(c)h grant no. 06999902 is gratefully acknowledged. We would also like to express our gratitude to Yunus Aksoy, Jan Behringer, Domenico Delli Gatti, Anna Gebhard, Bettina Gregg, Gabriel Lozada, Özlem Onaran, Benjamin Tippet, Jan Weber, Rafael Wildauer and the participants of the research colloquium at the University of Utah as well as the 12th Annual PhD Student Conference of the Post-Keynesian Economics Society, especially to our designated discussant Caverzasi Eugenio, the participants of the 25th conference of the Foundation for Macroeconomics and Macroeconomic Policies in Berlin, of the First Philosophy and Economics Conference in Vienna, of the Workshop on Economics with Heterogeneous Interacting Agents 2021, of the Networks in a Globalized Work Conference 2022 and of the Annual Conference of the European Association for Evolutionary Political Economy in Naples for many helpful comments and suggestions that greatly improved our paper and sharpened our focus. We are especially honoured to have received the Herbert Simon Prize 2022 for this paper by the European Association for Evolutionary Political Economy that was an immense help in gaining exposure and additional valuable feedback for our approach. We also thank one anonymous reviewer and the editor Morris Altman for their comments and excellent peer-review. Finally, we also thank Marton Gera, Carsten Källner and Marius Theiler for their able research assistance. All remaining errors are, of course, ours.

The nexus between debt and inequality has attracted considerable scholarly attention in the wake of the global financial crisis. One prominent candidate to explain the striking co-evolution of income inequality and private debt in this period has been the theory of upward-looking consumption externalities leading to expenditure cascades. We propose a parsimonious model of upward-looking consumption at the micro-level mediated by perception networks with empirically plausible topologies. This allows us to make sense of the ambiguous empirical literature on the relevance of this channel. Our model, based purely on *current* income, replicates the major stylised facts regarding micro consumption behaviour and is thus observationally equivalent to the workhorse permanent income hypothesis, without facing its dual problem of 'excess smoothness' and 'excess sensitivity'. We also demonstrate that the network topology and segregation has a significant effect on consumption patterns which has so far been neglected.

Keywords: Agent-Based Computational Economics, Behavioural Economics, Conspicuous Consumption, Inequality, Relative Income Hypothesis, Positional Goods, Emergence

1 Introduction

When the destabilising effects of both increasing economic inequality and excessive private debt became apparent in the wake of the global financial crisis, the nexus between the two has also increasingly attracted scholarly interest. One prominent candidate to explain the striking co-evolution of inequality and private debt in this period has been Veblenian conspicuous consumption resulting from the increase of top-end income inequality (Van Treeck, 2014). Veblen famously conjectured that "[...] our standard of decency in expenditures, [...] is set by the usage of those next above us in reputability; until, in this way, especially in any community where class distinctions are somewhat vague, all canons of reputability and decency, and all standards of consumption, are traced back by insensible gradations to the usages and habits of thoughts of the highest social and pecuniary class" (Veblen, 1899, p. 72). Following this line of argument, the increases in the top relative incomes triggered debt-financed conspicuous consumption of the poorer segment of the population and thus contributed to the financial fragility that ultimately led to the global financial crisis. While not being necessarily intended by anyone, the individual consumption behaviours thus collectively lead to a consumption "rat race" (Behringer et al., 2023). The income stratum directly below the super rich catch up with the increased consumption of the top income earners to preserve their relative position which in turn leads the income stratum below them to increase their consumption for the same reason. At the end of this collective process, aggregate and individual consumption has unanimously increased with their relative position in terms of expenditures remaining unchanged, though.[1]

While several microeconomic studies testify to the relevance of this upward-looking consumption behaviour (Heffetz, 2011; Bertrand and Morse, 2016; Bricker et al., 2020), the macroeconomic evidence on this channel has been more limited and ambiguous (Wildauer and Stockhammer, 2018; Bofinger and Scheuermeyer, 2019). In this paper, we argue that accounting for the meso-level institutional context that determines who compares their consumption to whom can explain the apparent mismatch between micro-behaviour and macro-aggregates. To integrate this institutional context, we build on two hitherto neglected mechanisms in Veblen's argument, summarised in the opening quote: First, social comparisons are local, i.e. pertain to those " next to us" and second, who is next to us is determined by the degree of "class distinction" that we interpret as network segregation. We attempt to capture both mechanisms in a parsimonious model that combines consumption emulation and reference groups determined by Veblen's degree of class distinction. Reference consumption activities to be emulated must be *observable* which we formalise by a plausible perception network (Schulz et al., 2022). Up to our knowledge, our approach is the first to make the reference group to which conspicuous consumption

[1] We thank an anonymous reviewer for pointing us to this argument and terminology.

relates explicit at a granular level. The sociological literature on everyday interaction determining reference coonsumption indicates that empirical social networks are far from random and consequently that human interaction is strongly structured. The major contribution of this paper thus is to show that the impact of empirically plausible networks, most notably networks featuring homophily or the tendency to link to individuals that are similar to oneself (Lazarsfeld et al., 1954), is non-negligible and might help explain the mismatch between micro- and macro-level findings. As predicted by Veblenian conspicuous consumption, increases in income inequality trigger higher consumption levels in our model and thus decrease aggregate savings. However, endogenously evolving perception networks due to homophily might mitigate this effect because an increase in inequality also increases segregation within the network. For high degrees of homophily, the effect vanishes, potentially explaining the contradictory empirical findings in the micro and macro literature.

On a methodological level, we outline a synthesis model of Veblenian conspicuous consumption, also drawing on various strands of heterodox consumption theory before and after Veblen. By highlighting these connections and compatibilities, we hope to contribute towards an alternative behavioural theory of consumption that is grounded in psychological and sociological theory and might complement conventional consumer theory based on constrained optimisation. The argument that intersubjective comparisons can be powerful personal motivators traces back to at least Adam Smith's (1776) *Wealth of Nations.* He famously argued that even something as simple as wearing a linen shirt might be a "necessary of life in England" (Smith, 1776, vol. ii, p. 352) at the time of his writing to signal that the wearer conforms to custom. Status signalling has been a central tenet of theories of consumption behaviour ever since. However, while most those theories focus on luxury goods, Smith (1776) reminds us that status goods are essentially defined by custom, and that their signalling is not an inherent feature of them. Therefore, status consumption is not confined to a specific social class that can afford expensive luxury goods but is a human tendency that might apply to all individuals across the whole income distribution. Yet, not all consumption is status consumption: Keynes (1972) divided goods into two types: Non-positional goods relating to material needs and status goods relating to needs "which are relative in the sense that we feel them only if their satisfaction lifts us above, makes us superior to, our fellows" (Keynes, 1972, p. 326).[2] Psychological evidence corroborates the conjecture that status consumption is upward-looking (Frank, 1985). Our model includes both the Smithian and Keynesian insight and lets *all* consumers have a propensity for *upward-looking* status consumption, with the whole population consuming a basket of both status and non-status goods.

In the Keynesian tradition, Duesenberry (1949) explicitly formalised a consumption function in which agents react to others' spending to "keep up with the Joneses". Despite its empirical success in replicating the known stylised facts regarding consumption expenditures at his time, the theory was quickly superseded by the now dominating permanent income hypothesis of Friedman (1957), likely due to Duesenberry's emphasis on sociological factors and behavioural interdependencies that are hard to reconcile with the general equilibrium paradigm (Mason, 2000). Most recently, Frank et al.'s (2014) concept of "expenditure cascades" combines aspects of the formal apparatus of Duesenberry (1949) with the Veblenian notion that expenditures by the poor can be partially traced back to the conspicuous consumption of the super rich: Given upward-looking consumption behaviour, uneven income gains that are skewed to the top can lead to cascading expenditures if the poorer segments of the population try to emulate the consumption behaviour of the richer strata (Frank et al., 2014). Our model builds on a similar formalism as in Duesenberry (1949) and Frank et al. (2014) and generates expenditure cascades in a perception network that captures social hierarchy. This parsimonious consumption rule based only on *current* income replicates all the major stylised facts regarding micro consumption behaviour. It is thus at least observationally equivalent to the workhorse permanent income hypothesis,

[2]The terminology of "positional goods" is due to Fred Hirsch (1976), not Keynes, though.

without facing its dual problems of 'excess smoothness' and 'excess sensitivity' (Meghir, 2004).

The remainder of this paper is organised as follows: Section 2 situates our model within the pertinent literature and shows how both sociological and economic theory inform our modelling choices. Section 3 introduces our formal model regarding consumption behaviour and the underlying social network. Section 4 presents our simulation results. Within the section, we demonstrate both that the emergent expenditure distributions are consistent with the relevant stylised facts and that the endogenous adjustment of the social networks can generate both nonlinearities and a large variety of point elasticities, underlining the potential relevance of the segregation channel. The final section 5 concludes and discusses possible avenues for further research, particularly possible empirical applications and modelling extensions to include our partial model of consumption decisions within a full-fledged macro framework.

2 Related Literature

The advent of Friedman's (1957) permanent income hypothesis was a paradigm change in the theory of consumption behaviour and shifted the focus from explaining consumption based on current income to explaining it with the concept of permanent income composed of unobservable stochastic income shocks. The subsequent literature has focused mainly on refining Friedman's initial hypothesis by including expectations, precautionary savings motives or liquidity constraints to bring it closer to the data (Palley, 2010; Meghir, 2004, for surveys). However, the dual problem of 'excess smoothness' and 'excess sensitivity' for calibrated consumption models of permanent income hints at the shortcomings of the permanent income hypothesis (Meghir, 2004), that is, consumption reacting too strongly (weakly) to (un)anticipated income shocks.[3] Moreover, the workhorse consumption model

[3] The commonly employed habit persistence extension to remedy these problems yields coefficients of habit formation that are typically deemed much too high to be plausible (Pesaran, 2015, p. 887 – 892, and the references therein, for a more detailed discussion).

built on Euler equations as famously introduced by Hall (1978) builds on rational expectations that are very demanding as an assumption. Empirically, Pesaran and Weale (2006) find little evidence for the rational expectations hypothesis. Theoretically, Hendry and Mizon (2014) show that a violation of the usual regularity conditions in stochastic processes like non-stationarity or structural breaks implies that rational forecasts conditioned on past performance are suboptimal, essentially since the application of the law of iterated expectations is hindered. Thus, it seems highly unlikely that the rational expectations benchmark indeed holds in the real world, especially if empirical income processes are very noisy.

We therefore abstract from any intertemporal expectation formation or postulating liquidity constraints and show how the emergent expenditure distributions relating to *current* income only can give rise to the relevant empirical stylised facts of consumption behaviour purely from human interaction within plausible social networks. Nonetheless, expectations regarding future income streams are potentially relevant for consumption behaviour, and our model merely demonstrates that the empirical findings on consumption expenditures both on the micro- and on the macro-level do not necessarily imply intertemporal optimisation over (largely unobservable) future income streams, as is typically argued in the literature (Battistin et al., 2009). In contrast to these demanding assumptions on (essentially unbiased or model-consistent) individual expectations, we make the perceptions of each agent fully explicit without any kind of expectation formation on future income streams. Thereby, we bring two strands of literature together to make sense of a third one: Namely, we combine the rich extant literature on empirical social networks with the empirical evidence on upward-looking consumption externalities to provide a new explanation for the ambiguous empirical findings on the aggregate savings-inequality nexus.

The topology of empirical social networks exhibits salient and universal features that can serve as stylised facts to validate artificial graphs. These include the small-world property which holds across many different domains (Weeden and Cornwell, 2020), homophily in tie-formation, especially for economic class (Cepić and Tonković, 2020; Malacarne, 2017; Mayer and Puller, 2008) and their sparsity, especially when focusing on close ties (Mac Carron et al., 2016; De Giorgi et al., 2020). Schulz et al. (2022) present a model of homophilic tie formation that can generate these stylised facts of empirical graph topologies as well as being able to replicate all the relevant findings on positional self-assessment and perceptions of inequality. Since our argument builds on status consumption, which rests on the individuals' perception of consumption, we adopt this modelling framework as an algorithm for plausible social networks.

One can understand this as individuals trying to maximise their social capital (Annen, 2003) by picking those others as social contacts who are like them in certain respects because this minimises the transaction costs of social interaction (Akerlof, 1997). Such transaction costs may simply be travel time if the social contacts do not live nearby, effort necessary to understand each other's cultural and social background or their interests. However, as factors like residential area (Harting and Radi, 2020), education (Leo et al., 2016), ethnicity (Chandra, 2000), lifestyle (Virtanen et al., 2007), or even health (Krieger, 1992) correlate with income, the latter constitutes a decent proxy for the actual homophilic behavioural patterns. Moreover, there is also homophily in income surfacing in, e.g. mobile phone communication (Leo et al., 2016). Therefore, while people may not willingly choose their social contacts based on income proximity, their actual choices amount to social contacts as if selected homophilic in income.

The importance of interpersonal comparisons in consumption and income is well established for the microeconomic level. Among others, Bertrand and Morse (2016), Bricker et al. (2020) and Heffetz (2011) find significant effects of interpersonal comparisons for individual consumption in the US, Jinkins (2016) for the US and China, Quintana-Domeque and Wohlfart (2016) for the UK and Drechsel-Grau and Schmid (2014) for Germany. However, the choice of reference groups within these studies often lacks granularity, underlining the need for more plausible models of interaction in networks. The empirical psychological literature reviewed in chapter 2 of Frank (1985) only establishes that conspicuous consumption is generally *upward-looking*. In addition to empirical studies of reference consumption, studies for the US also document the relevance of inequality for private debt-buildup of the poor to enable consumption increases, indirectly corroborating the relative income hypothesis (Agarwal et al., 2020; Carr and Jayadev, 2015).

Given the apparently unambiguous results from more micro-level studies, it seems surprising that this view has not manifested itself in a consensus on the macro-level association of inequality and (private) savings. While Koo and Song (2016) report that aggregate saving rates increase with income inequality due to the rich saving more, Klein (2015), Behringer and Van Treeck (2018), Behringer and van Treeck (2019) and Petach and Tavani (2021) find evidence for a negative relationship between savings and inequality, in line with the expenditure cascades hypothesis. Most studies fail to find a significant effect in either direction, though (Cuaresma et al., 2018; Gu and Huang, 2014; Gu et al., 2015; Leigh and Posso, 2009; Wildauer and Stockhammer, 2018). Cuaresma et al. (2018), Gu and Huang (2014) and Gu et al. (2015) highlight the potential relevance of unobservable intermittent country characteristics shaping the relationship between savings decisions and inequality, which might in part explain why Bofinger and Scheuermeyer (2019) find a strongly non-monotonic relationship between both.[4] Furthermore, Wildauer and Stockhammer (2018) highlight the time-scale of adjustments and show that with the proper controls and examining a long-run relationship, the effect of inequality on savings seizes to be significant.

We conclude that a proper model of expenditure cascades has to make sense of the variability of effect size as a function of the considered time scale. Moreover, such a model should give a plausible account of the origin of cross-country and even interregional heterogeneity. To demonstrate the relevance of network structures for these phenomena, we use an otherwise purely upward-looking consumption rule inspired by Frank et al. (2014) that

[4]Bofinger and Scheuermeyer (2019) include country- and time-fixed effects into their estimation. These might, however, be insufficient to capture time-variant country and region idiosyncrasies, as our endogenously evolving social network might suggest.

should in and of itself give rise to an unambiguously negative relation between savings and inequality but can generate a large variety of effect sizes as a function of network topology. Finally, we require our model to be consistent with the major stylised facts on empirical expenditure distributions for validation at the micro-level.

The empirical literature has identified four stylised facts regarding consumption expenditures and their relation to (current) incomes: (i) individual average propensities to consume (APCs) tend to decrease in (current) income (Dynan et al., 2004; Fagiolo et al., 2019); (ii) population-level APCs remain roughly constant with respect to changes in total income (Kuznets, 1942); (iii) the distribution of consumption expenditures is more homogeneous than the distribution of current income (Krueger and Perri, 2006; Jappelli and Pistaferri, 2010) and (iv) the distribution of consumption expenditures is (at least as a first-order approximation) well fit by a log-normal distribution (Battistin et al., 2009; Brzozowski et al., 2010; Chakrabarti et al., 2018; Hohnisch et al., 2002; Fagiolo et al., 2010), while the distribution of current income is not (Drăgulescu and Yakovenko, 2001; Silva and Yakovenko, 2004; Tao et al., 2019).

3 Model

This section gives a content-oriented presentation; we provide a technical description following the ODD protocol as well as NetLogo and Julia implementations of the simulation model on GitHub (Mayerhoffer and Schulz, 2023). The model consists of two distinct phases running in sequential order, where phase one builds the environment and determines available information for the consumption procedure in the second phase:

1. Network generation

    - Agent initialisation and income allocation
    - Homophilic linkage

2. Individual perception and consumption

Figure 1 shows the mechanisms featured in the model. Below, you find more detailed descriptions on each step and Section 4 explains the results following from the setup. The basic concepts are as follows: Agents form links based on homophily, i.e. with others who are close to them in income: the resulting perception network is more segregated the higher the homophiliy and the greater the income inequality. Indiosyncratic consumption is a fixed proportion of income for all agents but status consumption depends on the consumption of one's reference agent, given by the highest consumption of others that one observes in their ego-network. Consequently, there are two counteracting mechanisms induced by an increase in income inequality: Given a network, increases in aggregate inequality also increase inequality within perception groups and thus decrease savings. Yet, higher inequality also means a more segregated network, making inequality less visible and increasing savings. As discussed below, limited perception has an immediate effect (who is one's reference agent) and a mediate one (who is the reference agent's reference agent). Overall, poorer agents have higher levels of status consumption and higher APCs than richer ones, rendering expenditure levels more homogeneous; finally status consumption also means that there is a lower level of saving for everyone but the richest individual(s) not affected by status considerations.

3.1 Network Generation

There are 1000 agents in the model; each agent draws their income from an exponential distribution with a mean of $\lambda = 1$. Such a distribution normalises the empirical observed (pre-tax or market) income distributions in various industrialised countries for the vast majority of individuals (Drăgulescu and Yakovenko, 2001; Silva and Yakovenko, 2004; Tao et al., 2019; Shaikh et al., 2020). Thus, one can understand the model population as constituting a representative sample of empirical populations of these countries. The upper tail of 1 to 5 % of the income distributions empirically follows a Pareto law (Silva and Yakovenko,

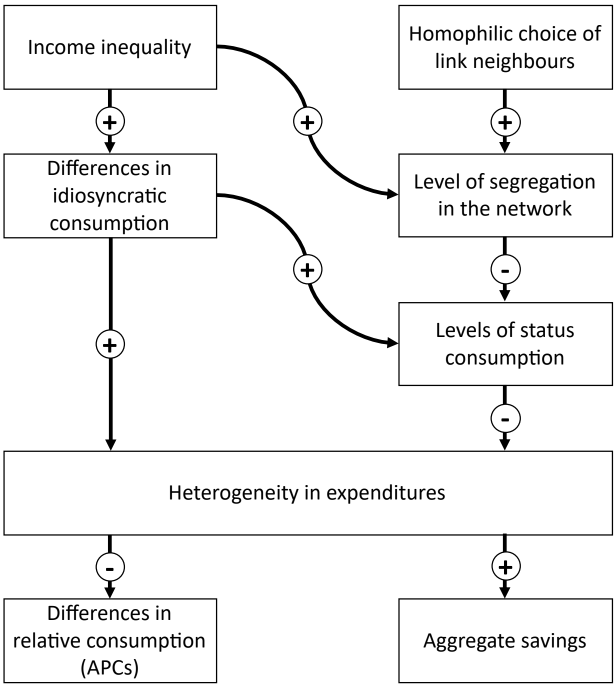

Figure 1: Mechanisms determining consumption and savings per model setup.

2004; Shaikh et al., 2020). We deliberately choose to exclude this small minority from our model since their population size would induce another degree of freedom in our model, and we want to demonstrate that segregation is indeed endogenous and not driven by differences in the actual income regime. Hence, we also use log-normal distributions for comparison between different inequality levels. While such log-normal distributions do not represent empirical findings, they allow for varying the inequality level through their dispersion parameter σ without changing the overall distribution shape. Hence, the present model employs the log-normal distribution solely to evaluate the relationship between actual income inequality, individual perceptions and consumption behaviour. Agents store their true income decile for evaluation purposes, too.

Each agent draws five other agents to link to. Like for real-world networks, links are therefore created *by agents*, not imposed on them. The number of five-link choices represents the closest layer of intense contacts identified in empirical studies (Zhou et al., 2005; Mac Carron et al., 2016)[5]. While this closest layer consists of differently related people for different individuals (Karlsson et al., 2005), what those individuals do, Wilcox and Stephen (2013) empirically identify this layer as the relevant reference group for social consumption. Moreover, we have carried out sensitivity analyses and found a larger number of links to have little impact on overall simulation results. 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 denotes the income of an agent, and $\rho \in \mathbb{R}_0^+$ denotes the homophily strength in income selection, externally set, and identical for all agents. $\rho = 0$ represents a random graph, and for an increasing positive value of ρ, an agent becomes ever more likely to pick link-neighbours with incomes closer to their own. [6] The link function's exponential character ensures that those with large income differences become unlikely picks even at low homophily strengths. Table 1 exemplifies the homophilic linkage for the median agent and expected incomes in the exponential distribution. For an extensive analytical discussion of this linkage behaviour, see Schulz et al. (2022).

While our linkage function represents a tractable reduced-form representation of the empirically well-established homophily, Falk and Knell (2004) demonstrate that endogenous reference groups can emerge from the social comparisons of agents aiming to self-improve and self-enhance. Similarly, our weighting function can be interpreted as agents selecting contacts as if minimising the income distance that they observe as a noisy signal. As discussed in Section 2, this assumption does not require individuals to consciously form

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

[6]One can understand the underlying linkage process as a generalisation of the preferential attachment process. The dimension on which the attachment is based can be any node characteristic: It is degree for the traditional preferential attachment and income difference in the present case but could for example also be income to model attachment to the richest. Furthermore, the parameter ρ enables the modeller to set the strength of the preference.

Income Rank	**Expected income**	**Income difference**	**Weight for** $\rho = 1$	**Weight for** $\rho = 4$
1	0.001	0.6921472	0.500500	0.06275
250	0.2876821	0.4054651	0.666667	0.19753
499	0.6911492	0.001998	0.998004	0.99204
500	0.6931472	0	-	-
501	0.6951492	0.002002	0.998000	0.99202
632	0.99967	0.30653	0.736000	0.29343
750	1.38629	0.6931472	0.500000	0.06250
1000	7.48547	6.7923237	0.001120	1.587e-12

Table 1: Exemplary linkage for the median agent (income rank 500).

their social ties based on income proximity. Instead, they may exhibit homophily in e.g. education level, ethnicity, or lifestyle; but thereby the individuals act as if there was explicit income homophily because the aforementioned factors correlate with income. As a consequence, segregation increases with income inequality, as is well documented in the literal sense of residential segregation (Chen et al., 2012; Tóth et al., 2021). However, links in our model extend beyond this literal geographic sense and essentially imply that consumption decisions are observable, be it due to common workplaces, family ties or any other form of linkage.

Formally, we assume that the utility of agent i connecting to agent j follows an additively separable utility function with agents receiving disutility proportionally to the absolute income distance to the income of agent j and a random error term, i.e.,

$$U_{ij} = -\rho|Y(i) - Y(j)| + \epsilon_{ij}. \tag{2}$$

When the distribution of ϵ_{ij} is identically and independently distributed following an extreme value distribution, we can express the choice probability of agent i as

$$p_{ij} = \frac{\exp[-\rho \cdot |Y(j) - Y(i)|]}{\sum_{k \in M \setminus i} \exp[-\rho \cdot |Y(k) - Y(i)|}, \tag{3}$$

with $M \setminus i$ as the set of all agents except i with size $N-1$ which is equivalent to the weights in eq. (1) translated into probabilities (Hoffman and Duncan, 1988). Our weighting function is thus an application of the discrete choice approach developed and popularised by Manski and McFadden (1981). The random utility model appears to us to be particularly appropriate here, since income is a rather salient characteristic to determine a good 'fit' within a social network but of course might depend on other characteristics that are not directly observable and thus modelled to be stochastic.[7] The homophily parameter $\rho \in (0, \infty)$ is then simply the intensity of choice parameter with $\rho \to \infty$ implying that i chooses agent j to link to with certainty who has minimal income distance, i.e., $p_{ij} \to 1$. In this sense, the weighting function in eq. (1) is plausibly microfounded and can now be considered the workhorse choice rule in behavioural macroeconomics (Franke and Westerhoff, 2017). Franke and Westerhoff (2017) also survey evidence from a several lab experiments in different macroeconomic contexts that discrete choice is indeed consistent with the data, while Anufriev et al. (2018) provide laboratory evidence for the discrete choice approach for financial markets.

Our model thus assumes that homophily solely results from the utility maximisation of agents, implying that all homophily is *choice* homophily. *Induced* homophily by shared workplaces or educational stratification, for example, per construction does not play a role in our model. We assume this to preserve tractability and not to counterfactually argue that induced homophily is empirically irrelevant (McPherson and Smith-Lovin, 1987; Chetty et al., 2022). For our simulation results, only the empirical segregation based on income levels is relevant since both types of homophily are captured by the parameter ρ. Nevertheless, while this distinction of different types of homophily is irrelevant for

[7]Note, however, that this choice rule implies the axiom of Independence of Irrelevant Alternatives (IIA) to hold for all alternatives and was in fact explicitly designed to do so (Manski and McFadden, 1981), i.e., the probability of choosing between j and k to be independent of the probability of choosing l. IIA might be considered a good first-order approximation to model homophilic choice, but especially within social networks, knowing one agent j might indeed increase the likelihood of knowing another agent l that is acquainted with j. It might thus prove interesting to extend and generalise the above rule for correlated choice to examine the effects on the network topology in further research.

the formal structure of the model, it might matter for policy: If homophily is mainly structural, it can also be influenced by targeting structural characteristics like housing subsidies to counter residential segregation. The individual preferences on which choice homophily are based are arguably much harder to alter for policy. Disentangling these two types of channels in our model thus constitutes a promising avenue for further research.[8]

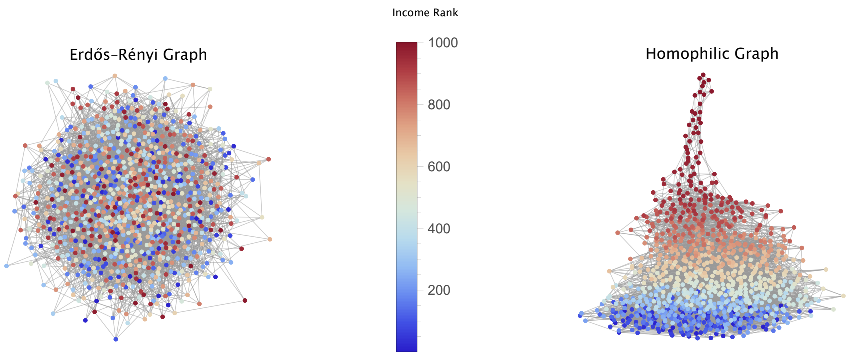

Figure 2: The figure illustrates the effect of income homophily on network segregation. The left panel corresponds to a fully random (Erdős–Rényi) graph with $\rho = 0$, while the right panel shows a homophilic graph with $\rho = 4$. While the random graph is expectedly completely unstructured, the homophilic graph is segregated according to income class. Graph visualisation by Mathematica's "Spring Electrical Embedding" routine.

Figure 2 contrasts a purely random graph with $\rho = 0$ with a homophilic graph initialised with $\rho = 4$.[9] As the figure clearly demonstrates, Veblen's class distinctions only exist for the homophilic graph on the right panel, while the random graph on the left panel exhibits no stratification at all. We thus argue that we can interpret the homophily parameter ρ as the degree of class distinction put forward by Veblen. For a growing ρ,

[8] We thank an anonymous reviewer for pointing us to this distinction and the different implications for both types of homophily.

[9] We thank an anonymous reviewer for prompting us to investigate the network topology more deeply.

agents increasingly tend to interact with others in the same social stratum. Regarding percpetion, the emergent networks from a homophilic graph-formation process generate networks that replicate the relevant stylised facts regarding both overall inequality (mis-)perceptions (Schulz et al., 2022) as well as perceptions of intergroup inequities, i.e., the gender and racial wage gaps (Mayerhoffer and Schulz, 2022a,b). Since perceptions are empirically based on these social samples, the validity in terms of perceptions these studies establish translates into external validity of the network we employ to model everyday interactions. The fact that the generated networks also exhibit real-world topological properties of social networks adds to the plausibility of the generating process (Schulz et al., 2022).

3.2 Individual Perception and Consumption

To avoid the problems related to the conventional approach to consumption Euler equations presented in Section 2, we base our model on current income and intersubjective comparisons. Moreover, this approach does not require an unrealistic degree of heterogeneity or additional assumptions to match the empirically observed heterogeneity in consumption propensities (Jappelli and Pistaferri, 2014).

In particular, the only heterogeneity we impose is the (pre-validated) income distribution. Agents consume maximising utility $U(\cdot)$ with Cobb-Douglas preferences and derive (linear) disutility the stronger their consumption in isolation given by $\tilde{C}(i)$ falls short of the highest consumption level $C(j|i)$ they observe within their ego network with intensity parameter $c > 0$. With this assumption, we follow the tradition of additive comparison utility in contrast to ratio comparison utility, where others' consumption enters as a ratio to own isolated consumption.[10] Apart from that the utility function is rather standard with agents deriving utility both from consumption with elasticity parameter w and saving with

[10] A prominent example of the additive approach is Akerlof (1997), while one example of ratio comparisons can be found in Carroll et al. (1997). The qualitative intuition should be equivalent in both cases, though.

elasticity parameter $(1-w)$ to capture intertemporal motives in the most parsimonious fashion. The choice variable is the average propensity to consume γ, i.e., the fraction of current income Y an agent consumes. The utility function of agent i thus reads

$$U(\gamma; Y(i), C(j|i), \tilde{C}(i), c, w) = (\gamma_i Y(i) - c \cdot (C(j|i) - \tilde{C}(i)))^w ((1-\gamma_i)Y(i))^{1-w} \tag{4}$$

The utility function for consumption in isolation, $\tilde{U}$ with choice variable $\tilde{\gamma}_i$ is given by eq. (**??**) for $c = 0$ and thus reads

$$\tilde{U}_i(\tilde{\gamma}_i; w, Y(i)) = (\tilde{\gamma}_i Y(i))^w (Y_i \cdot (1-\tilde{\gamma}_i))^{1-w}. \tag{5}$$

The FOC for optimal $\tilde{\gamma}$ reads

$$\tilde{\gamma}_i = w, \tag{6}$$

and thus optimal consumption without social interaction is given by

$$\tilde{C}(i) = wY(i), \tag{7}$$

i.e., the canonical result that consumption is a constant fraction of income. Substituting $\tilde{C}(i)$ from eq. (7) into eq. (5) yields

$$U(\gamma; Y(i), C(j|i), c, w) = (\gamma_i Y(i) - c \cdot (C(j|i) - wY(i)))^w ((1-\gamma)Y)^{1-w} \tag{8}$$

By the FOC we derive the optimal γ_i as

$$\gamma_i = \frac{wY(i) + (1-w)c(C(j|i) - wY_i)}{Y(i)}, \tag{9}$$

and consequently consumption by noting that $C(i) = \gamma_i Y(i)$ as

$$C(i) = wY(i) + (1 - w) \cdot c(C(j|i) - wY(i)). \tag{10}$$

Since we assume w and c to be constant for all agents, the consumption rule is entirely equivalent for all agents. This implies that the consumption of any agent i derived here from a utility maximising framework is a weighted sum of idiosyncratic and socially determined consumption with weights w and $(1 - w)$ (Frank et al., 2014).

We want to emphasise that the derivation by utility maximisation is not the only conceivable way to arrive at such a consumption rule. Equivalently, the rule can also be interpreted as resulting from behavioural heuristics, where w is the marginal propensity to consume out of current income and c is the propensity to consume out of the consumption differential to the individual indexed j without any reference to utility functions. In the terminology of evolutionary theories of consumption, the first term in (10) would capture the 'needs' of consumers with the latter term corresponding to the socially constructed 'wants' (Witt, 2001). In Keynesian jargon, these terms would instead be interpretable as capturing the intuition of the absolute and relative income hypothesis (Palley, 2010), respectively. The formal representation is in this sense neutral with respect to the subtleties of marginalist, evolutionary or Keynesian interpretations and can thus be incorporated in those different frameworks without much difficulty.

Overall, intertemporal motives only feature within the model insofar as agents derive utility directly from saving in anticipation of future consumption and hence leave no role regarding expectation formation regarding potential future income gains or losses. However, our approach is also consistent with the permanent income hypothesis's central intuition: One can also understand the interpersonal consumption comparison presented here as a proxy for intertemporal ones. Individuals might look at others they perceive to lead overall parallel lives as themselves and take the observed consumption as a predictor for what is possible for themselves. In that way, our model explains why life-cycle models replicate the four stylised facts but necessitate implausible behavioural assumptions to do so: Individuals act as if they smoothed over future incomes while actually observing some of these potential future incomes in others. Arguably, merely observing consumption decisions rather than forming explicit (and model-consistent) expectations over a stochastic income innovation process appears to be much more plausible.[11]

Given the model mechanism, the income levels $Y(j|i)$ on which the consumption of i (directly or indirectly) depends can be indexed by an ordering $\mathcal{R}_i = 0, 1, ..., d_i$ where d_i is the distance of the most distant income to i on which $C(i)$ depends. It follows that $Y(0|i) = Y(i)$, as the distance is then 0. By recursive substitution, we can rewrite $C(i)$ as

$$C(i) = \sum_{j=0}^{d_i} ((1-w)\cdot c)^j \cdot (1-(1-w)\cdot c)\cdot \; w \cdot Y(j|i) = \sum_{j=0}^{d_i} \beta_j Y(j|i). \tag{11}$$

The consumption of any individual i is thus directly dependent on their own income but depends also on all other incomes to which they are (directly or indirectly) connected. Since $((1-w)\cdot c)$ is, per assumption, strictly below unity, the weight of a given distant income level $((1-w)\cdot c)^j$ decays (exponentially) in the graph distance $j = 0, 1, ..., d_i$ to the relevant agent. While more distant incomes are thus generally higher due to social comparisons being upward-looking, they are also given less weight. This implies that $C(i)$ is a linear combination of an exponentially distributed random variable Y with varying weights and number of terms and thus, a so-called hypoexponential mixture (Li and Li, 2019). As we discuss in more detail in Section 4, this directly implies two of the four stylised facts, namely, the approximate log-normality of expenditure distributions and the fact that they are robustly more homogeneous than income distributions.

[11] We thank Eugenio Caverzasi for pointing us to this alternative interpretation.

4 Results

The results we present are twofold: Firstly, we show that our parsimonious model setup embedded within realistic social networks can replicate the known stylised facts of empirical expenditure distributions at the micro-level. On the one hand, these results testify to the validity of our model assumptions. On the other hand, they also serve as proof of concept to show that a static setting may replicate empirical expenditure patterns without resorting to unobservable stochasticity, strong assumptions about credit markets and individual rationality or assuming any heterogeneity in individual consumption functions ex-ante. Secondly, we spell out the implications for savings and inequality at the macro level. In particular, we show that our model can generate a large variety of different elasticities of savings with respect to inequality, in line with the highly ambiguous empirical literature.

4.1 Micro Level Patterns

As laid out in Section 2, we consider four stylised facts for validation with varying levels of granularity. Stylised facts (i) to (iii) describe qualitative features of expenditures that are, in principle, consistent with many different functional forms describing the distribution of expenditures. The fourth stylised fact is more stringent and imposes log-normality as the functional form. Even though log-normality is a more demanding requirement, it does not itself imply stylised facts (i) to (iii), which we, therefore, consider separately for full validation (Fagiolo et al., 2019). The first stylised fact concerns the decrease of individual APCs in income levels. Indeed, our model replicates this finding for all levels of social consumption without imposing different consumption propensities ad hoc. This follows directly from social consumption being purely upward-looking. If one is richer than all their link-neighbours, they do not need to catch up to anyone's higher consumption, and hence their consumption is entirely idiosyncratic. However, social consumptfion accumulates as it is passed down within the network from richer to poorer individuals. Thus, poorer indi-

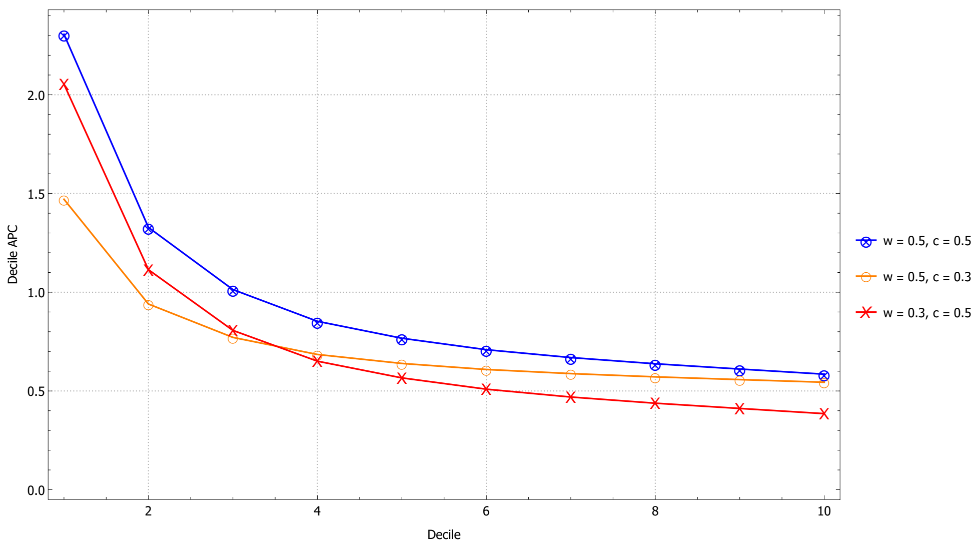

Figure 3: The figure plots the APCs per income decile for different parametrisations and with homophily strength $\rho = 4$ for a single simulation run each. Lines are visual aids only.

viduals tend to exhibit higher social consumption above their idiosyncratic consumption, increasing their personal APC. Figure 3 shows the APC schedules per decile for different parametrisations and demonstrates that decile APCs are indeed declining, in line with the empirical evidence. Note that this result is endogenous to our model and not imposed *ex ante*.

As Figure 3 also shows, though, the precise form of this decay depends on the specific parametrisation. The w parameter determines the weight of idiosyncratic consumption. Since idiosyncratic consumption is relatively more important for the richest, w is also most relevant for the level of the total consumption of the richest decile, as is apparent from the fact that the blue and orange APC schedules with equal w but different c in Figure 3 almost coincide for the richest decile, whereas the red and orange schedules with equal c but different w do not. In contrast to that, the c parameter captures the strength of the 'catching-up' behaviour, which manifests itself in the steepness of the curve, as is also readily visible from comparing the orange and red APC schedules.[12] Regardless of

[12]We chose these particular parametrisations, since, for $\rho = 4$, they appear to capture the behavior of empirical APCs reported in Clementi and Gianmoena (2017) reasonably well. Decreasing ρ tends to increase the reported effects since the income difference to the richest observed individual within

the parametrisation leading to different functional forms of the APC schedules, they are, however, all declining in income decile, replicating stylised fact (i).

The most prominent fact regarding empirical expenditure distributions was the finding by Kuznets (1942) that aggregate APCs stay approximately constant for changes in aggregate income. Indeed, this was one of the major factors contributing to the paradigm change from the absolute to the permanent income hypothesis, as the absolute income hypothesis in its affine-linear variant exhibits decaying aggregate APCs in total income (Palley, 2010). The invariance of aggregate APCs C/Y from changes in Y follows trivially from the recursive solution to $C(i)$ in eq. (11). Assume a proportional change to Y by $\omega \in \mathbb{R}^+$, such that $\tilde{Y}(i) = \omega Y(i) \quad \forall i$ and $\tilde{Y} = \omega Y$. Thus, ω is independent of $Y(i)$ and does not change orderings $\mathcal{R}_i$. Consequently, since ω cancels out, total consumption is given by the following:

$$\tilde{C} = \sum_{i=1}^{n} \sum_{j=0}^{d_i} \beta_j \omega Y(j|i) = \omega \sum_{i=1}^{n} \sum_{j=0}^{d_i} \beta_j Y(j|i) \text{ and therefore} \tag{12}$$

$$\frac{\tilde{C}}{\tilde{Y}} = \frac{\omega \sum_{i=1}^{n} \sum_{j=0}^{d_i} \beta_j Y(j|i)}{\omega \sum_{i=1}^{n} Y(i)} = \frac{C}{Y}. \tag{13}$$

As a result, aggregate APCs are invariant to any positive proportional change to income, i.e., distribution-preserving changes. Whenever inequality by one of the usual measures increases, aggregate APCs increase, too, as we show in the subsection below.

As a third stylised fact, we considered that the distribution of consumption expenditures is more homogeneous than the income distribution. In our specification, this is necessarily also true for $w < 1$,[13] since by the recursive solution in eq. (11), expenditure levels are a weighted sum of income levels, themselves following an exponential distribution, with varying weights and number of terms. This characteristic implies that the distribution of

a perception set tends to increase with decreasing ρ by assumption, leading to stronger 'catching-up'. Simulation results verifying this finding are available upon request.

[13]For $w = 1$, we would simply recover a (rescaled) version of the exponential income distribution for expenditures, and the recursive solution in eq. (11) would be undefined.

C is *hypoexponential*, with a coefficient of variation strictly smaller than unity, while the exponentially distributed income levels exhibit a coefficient of variation (asymptotically) equal to unity (Li and Li, 2019). Intuitively, since social consumption accumulates in a cascade down the income distribution, richer individuals are relatively unaffected by status concerns, while the cumulative effect on poor individuals is much higher. Thus, increasing the level of status consumption by decreasing w or increasing c tends to equalise the expenditure distribution since the relatively higher status consumption of the poorer lets their consumption approach the expenditures of the rich that are not so much affected by status concerns.

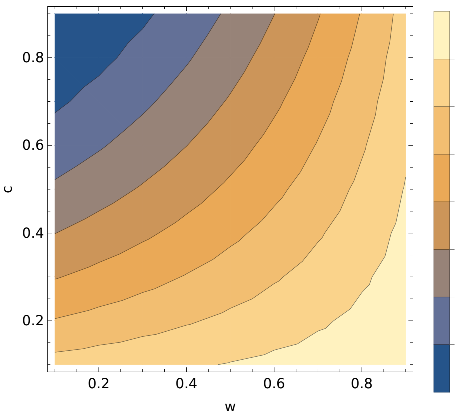

Figure 4: Ratio of the coefficients of variation for the expenditure and income distribution for the whole parameter space of w and c and with $\rho = 0.5$, averaged over 100 simulation runs.

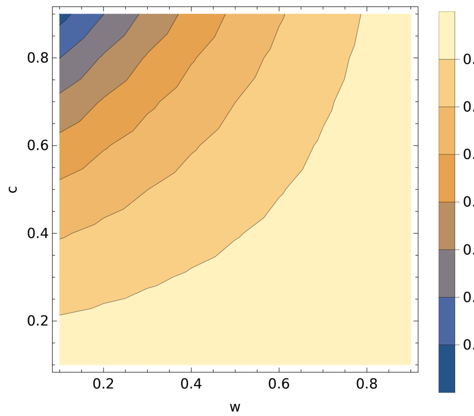

Figure 5: Ratio of the coefficients of variation for the expenditure and income distribution for the whole parameter space of w and c and with $\rho = 4$, averaged over 100 simulation runs.

This is also what we find for our numerical simulations, here for 100 runs each, as we show in Figures 4 and 5 for the parameter combinations of w and c. For all cases, the coefficient of variation is strictly below the one for the income distribution, replicating stylised fact (iii). However, the equalising tendency is much less pronounced for high levels of homophily, as can be seen in Figure 5 because the homophilic link formation tends to decrease the maximum income level each individual observes, as individuals only interact with others within their respective income stratum. Hence, homophily tends to hinder the transmission of social consumption.

Finally, we also find that our specification can map exponential income distributions to log-normal expenditure distributions, even though both come from completely different distributional families. The reason behind this finding is simply that the hypoexponential mixture characterising the expenditures is well approximated by a log-normal if it exhibits sufficient skewness, but the underlying rate parameters λ of the exponential are not too heterogeneous. A representative example is shown in Figures 6 and 7.

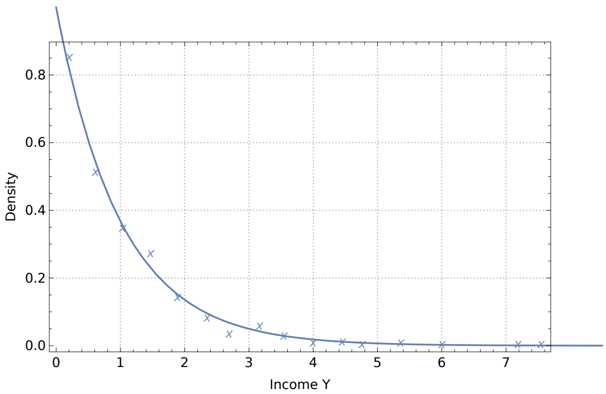

Figure 6: Empirical Density of Personal Incomes $Y(i)$.

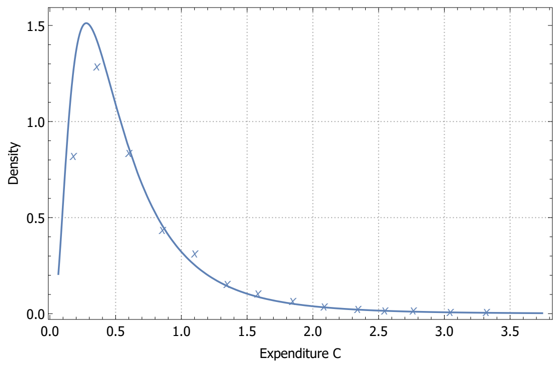

Figure 7: Empirical Density of Personal Expenditures $C(i)$ for a single simulation run with $w = 0.5$, $c = 0.3$ and $\rho = 1$.

Per assumption, the pre-validated income distribution follows an exponential distribution with rate parameter $\lambda = 1$. However, the expenditure distribution is extremely well fit by a log-normal distribution, replicating stylised fact (iv). Note that this represents an entirely new generating mechanism for log-normal distributions and does not rely on any unobservable income innovation process or model-consistent expectations as are typically used in mainstream macro models. This is a relevant finding in and of itself since the log-normality of consumption expenditures is typically taken as evidence that consumption cannot be a function of *current* income and thus needs to depend on *permanent* income subject to a stochastically multiplicative growth process generating the log-normal functional form (Battistin et al., 2009). [14] Our model exercise demonstrates that this is not necessarily true. A parsimonious model based on current income and simple behavioural rules of thumb can simultaneously replicate this stylised fact without facing the theoretical and empirical shortcomings of models in the rational expectations tradition.

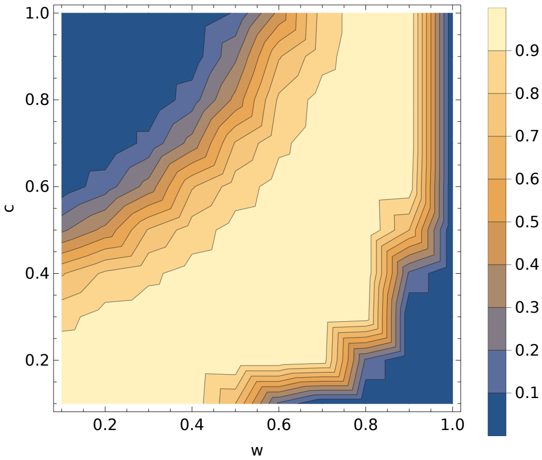

Figure 8: Fraction of expenditure distributions for which a standard Kolmogorov-Smirnov test cannot reject the assumption of log-normality at a 5% significance level for $\rho = 0.5$ with 100 iterations.

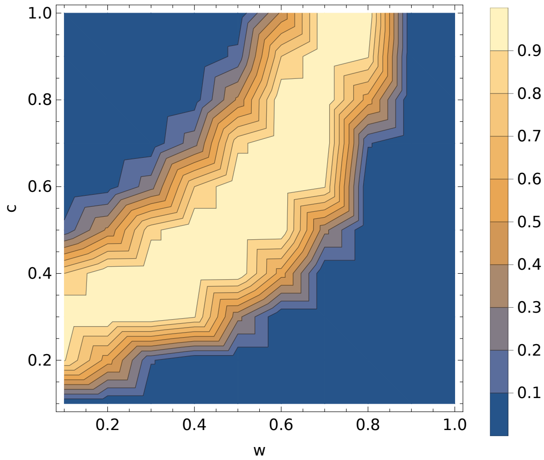

Figure 9: Fraction of expenditure distributions for which a standard Kolmogorov-Smirnov test cannot reject the assumption of log-normality at a 5% significance level for $\rho = 4$ with 100 iterations.

Log-normality holds for large parts of the parameter space, as we show in Figures 8 and 9. The approximation only breaks down for very high levels of idiosyncratic consumption and w approaching unity, where we recover a rescaled version of the initial exponential distribution and very high levels of social consumption with low w and high c, where the expenditure distribution becomes too symmetric to account for the skewness of the log-normal. For high levels of homophily ρ, the parameter space for log-normality becomes much narrower due to the emergent segregation that weakens the transmission of status

[14]This "law of proportionate effect" is perhaps the most prominent generating mechanism for log-normal distributions (Mitzenmacher, 2004). It lets permanent incomes grow randomly and multiplicatively but independently of scale, i.e., income growth rates are independent of income levels (Schulz and Milaković, 2023). While top incomes are dominated by capital incomes for which competitive financial markets might indeed imply multiplicative growth with equalising rates of return, it is unclear why this process should also govern the growth dynamics for the lower income strata. In fact, the distributional evidence rather points to additive growth for the vast majority of wage earners (Shaikh et al., 2020). In this sense, there might even be reason to discard the rather implausible assumption of the law of proportionate effect for the whole income distribution in favour of our explanatory attempt based on social emulation in networks.

consumption. In light of these findings, it seems entirely unsurprising that many of the empirical studies we reviewed in Section 2 find significant deviations from log-normality in expenditure distributions, which one can understand as arising from variation in consumption baskets' salience of expenditures as well as differences in social segregation. Both the heterogeneity in the salience of consumption goods (Solnick and Hemenway, 2005; Heffetz, 2011, for a large scale study) and the variation in (e.g., spatial) segregation (Tóth et al., 2021) are well documented in the literature, in agreement with the varying results on the log-normality of expenditures across studies. Perhaps more surprisingly, we find that our stylised consumption function embedded within realistic social networks can replicate the relevant empirical phenomena in the literature on empirical expenditure patterns, testifying our modelling approach's validity.

4.2 Inequality and Aggregate Savings

So far, the model assumed exponentially distributed incomes with a constant theoretical Gini coefficient of $G = 0.5$, irrespective of its particular parametrisation. To examine the impact of *changing* inequality on consumption and savings, we therefore now assume log-normality of individual incomes to generate variation in inequality through changing the dispersion parameter σ of the log-normal income distribution by which the model is initialised. We examine different degrees of homophily and let social networks adjust endogenously to input inequality, thereby highlighting two different findings: Firstly, there exist strong non-linearities in the relation of inequality and savings whenever perception networks mediate actual inequality. Secondly, the time scale of adjustment might be important. We only provide comparative statics here and report results after all adjustments regarding perceptions have taken place but hint at possible mechanisms in the time domain as well.

Figure 10 depicts aggregate saving rates as a function of income inequality, measured by the Gini coefficient, and for varying levels of income homophily. As expected within a

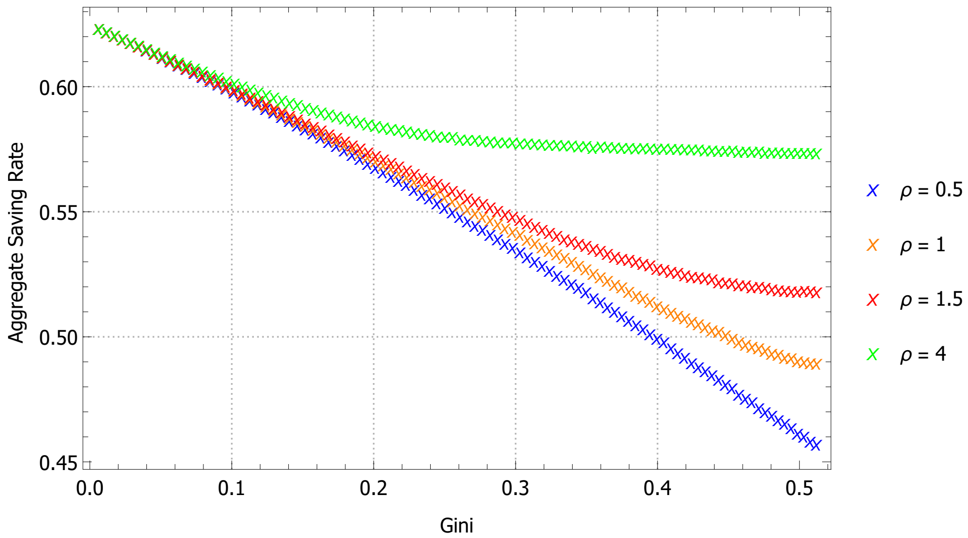

Figure 10: Aggregate saving rate (aggregate savings over aggregate income) as a function of inequality, measured by the Gini coefficient with $w = c = 0.5$ and $b = 0.75$, averaged over 100 simulation runs.

model of purely upward-looking consumption, savings unanimously decrease in inequality. Consequently, our approach replicates the main finding of the canonical model of 'expenditure cascades' by Frank et al. (2014) within plausible and endogenously evolving social networks. The reason for that are the two counteracting effects of increased income inequality shown in Figure 1.

For (implausibly) low degrees of inequality, the four trajectories for $\rho = 0.5, 1, 1.5$ and 4 almost appear to coincide, as the segregation-effect is negligible here, while the within-effect is fully passed through to savings. From a Gini coefficient $G = 0.1$ onwards and now within the empirically relevant range (Bofinger and Scheuermeyer, 2019), the trajectories begin to deviate, with still strongly negative effects for the low-homophily regimes but vanishing savings elasticities within the highly homophilic regime. This is apparent from the decaying schedules for $\rho = 0.5, 1$ and 1.5 in contrast to the plateau for $\rho = 4$. Methodologically, this finding thus points to emergent behaviour at the macro level that is irreducible to the individual micro-level behaviours: With completely identical behavioural rules at the micro-level, one can generate vastly different aggregate savings elasticities at the macro level. Reminiscent to Kirman's (1992) foundational demonstration, the properties even of identical micro agents do not translate to equivalent properties of a 'representative agent' whatsoever, as the nonlinear macro behaviour in Figure 10 in contrast to the identically linear micro consumption rule shows.

This methodological finding is not only of methodological interest but also bears material implications: Most importantly, it might in part explain why the micro literature appears to confirm the relevance of the conspicuous consumption channel, while the macro literature remains at best ambiguous on this issue. It might very well be the case that all individuals engage in conspicuous consumption without this being observable using only macro-level aggregates, whenever there is an endogenous selection of reference groups. Indeed, the scarce extant empirical literature suggests both that inequality increases (geographical) segregation (Tóth et al., 2021) and that (geographical) segregation decreases conspicuous consumption (Bertrand and Morse, 2016). Policies aimed at decreasing segregation might thus carry unintended consequences in simultaneously increasing conspicuous and potentially wasteful consumption. This adds an additional layer to problems in overcoming segregation pointed out first by Sakoda (1949, 1971) and much later by Schelling (1971) (see Hegselmann (2017) for an extensive discussion of the history of the checkerboard model).

Finally, our results are also a cautionary tale against the pitfalls of pooling macro data for analysis without acknowledging the potential time- and cross-sectional variation of inequality effects on segregation and perceptions. Indeed, it appears likely that the timescale of the homophily mechanism outlined in section 3 crucially depends on the mobility of actors that enables homophilic choice or induced homophily in the first place (Thomas, 2019, for a recent empirical study). With low mobility, the transitory adjustment to the new stationary state we depict in our comparative statics exercise in Figure 10 might take sufficiently long to be of macroeconomic relevance. In contrast, adjustment might be very fast for high degrees of mobility. The significant short-term elasticities, e.g. in Petach and

Tavani (2021) can therefore be reconciled with the insignificant long-term effects Wildauer and Stockhammer (2018) find by considering heterogeneous mobility in contrast to ad-hoc assumptions regarding behavioural parameters.

Furthermore, the degree of homophily itself appears to be varying across and within countries (Cepić and Tonković, 2020). Consequently, pooled estimates might obscure unobserved cross-country heterogeneity both in adjustment speed and strength, which might help to explain the non-monotonicity Bofinger and Scheuermeyer (2019) find. Therefore, our model also calls for better data to capture these effects, either by controlling for, e.g., residential or occupational segregation or by directly including measures of perceived incomes. Unfortunately, these attempts are most of the time hindered by data that is much too coarse-grained for proper estimation, with the geographical location only being available at the state or county level (Bertrand and Morse, 2016) and perceived incomes being reconstructed from data on a rather crude ordinal scale (Choi, 2019) or based merely on shared workplaces (De Giorgi et al., 2020). Merging the relationship network, including friends and family besides coworkers with the theoretically relevant socioeconomic data is rarely achieved in practice but would be an important step to further bridge social network and consumption theory (De Paula, 2017).

5 Discussion

Our goal was to provide a tractable evolutionary alternative to orthodox theories of consumption based on Euler equations and intertemporal optimisation. In our model we show that this can be achieved by much more parsimonous assumptions based on the distinction between 'needs' and 'wants' (Witt, 2001). With respect to the major stylised facts, the simple rule proposed here and the one induced by intertemporal optimisation appear to be observationally equivalent and even microdata on expenditure distributions might be unable to discriminate between the two. Theory, even in the form of carefully calibrated macro models, might be *underdetermined* in this regard (Quine, 1975). Intertemporal motives might matter for consumption, but the log-normality of expenditures is perhaps not an as decisive argument for consumption smoothing as it is often perceived to be (Battistin et al., 2009). Methodologically, intersubjective rather than intertemporal consumption rules might also be appealing since they build on much less restrictive assumptions. Intertemporal optimisation essentially implies both quantifiable income risk rather than fundamental uncertainty in the form of a well-defined probability distribution over income shocks and demands that this distribution is learnable or transparent for consumers (Menz, 2010). By contrast, social consumption rests on cross-sectional observables made explicit by the perception network and therefore requires agents to be mere observers rather than econometricians which strikes us as a more plausible portrayal of human behaviour. At the same time, our micro-level mechanisms are consistent with the permanent income hypothesis, too. By acting on what they see in others they perceive to be similar to themselves, individuals behave as if they were acting on expectations about their own lives. Hence, using social networks relying on interpersonal consumption comparisons might be attractive for future research as a proxy for intertemporal considerations and its own sake because conspicuous consumption itself is an empirically well-grounded phenomenon.

Our second major result concerns the relationship of individual consumption with aggregate consumption levels. We find that aggregate consumption might react far less elastic to changes in income than each individual consumption rule considered separately would indicate, even when all individuals unanimously follow exactly the same rule. This apparent contradiction follows from the endogenous selection of consumption reference groups. Apart from increasing aggregate consumption, inequality also increases segregation and therefore mitigates its initial consumption-enhancing effect. Naturally, this new channel sheds light on recent puzzles concerning the apparent divide between the micro and macro strands of the empirical literature on this matter. In particular, we explicate the potential micro origins of the institutional channel that Behringer and van Treeck (2022) and Ascione and Schnetzer (2022) find to be of crucial relevance to explain the empirical cross-country heterogeneity as well as time variation: In addition to the ease of access to credit or consumerist social norms, also the social network of everyday interactions might be an important determinant for relative importance of expenditure cascades. The growing literature on ‘growth models’ within comparative political economy might thus also benefit from considering institutional differences in social segregation (Behringer and van Treeck, 2019).

Moreover, our results point to a relevant trade-off for policy-makers: According to our model results, policies aimed at reducing occupational or residential segregation might also increase conspicuous consumption that might be environmentally wasteful (Howarth, 1996) and contributes to the destabilising build-up of private debt (Van Treeck, 2014). To offset these effects, our framework suggests complementing such policies of integration by appropriate taxation or behavioural nudges that are disincentivising status consumption, or, in the language of our model, decrease the catching-up parameter c.

The choice in favour of a unanimous consumption rule based on comparisons is a deliberate one for two reasons: Firstly, our parsimonious model’s internal validation was unproblematic as we could monitor the whole parameter space of all exogenous variables and did not detect any inexplicable irregularities. Secondly, we keep individual behaviour intentionally simple to show that the *interaction* of agents is sufficient to generate empirical heterogeneity in micro observables and a large variety of macro-level effects. Even in its parsimonious form, the present model provides an epistemically possible how-possibly explanation Grüne-Yanoff and Verreault-Julien (2021) whose explanans can be true in the real world and whose replicated stylised empirical facts is in “qualitative agreement with empirical macrostructures” (Fagiolo et al., 2019, p. 771). Finally, both the principle of homophilic choice and social consumption effects are well established in the extant empirical literature, further testifying to the resemblance (Mäki, 2009) between our model and reality. Therefore, our model reconciles technical verification with both input and descriptive output validation which Gräbner (2018) considers a rare feature of models. We are confident that the how-possibly explanation given by our model is preferable to ad-hoc choices and much more general or could at least be straightforwardly extended in this direction.

Nonetheless, the minimal nature of our model points to several limitations and extensions. Most notably, in its current bare-bones form, the model cannot study any issues of liquidity constraints that are frequently found to be of crucial relevance, especially for low-income households (Dynan et al., 2004) or spill-overs of private consumption to other sectors that would require the model incorporate general equilibrium effects. Generally, links are modelled abstractly and incorporate various types of personal or professional connection layers as proximity of incomes. Future work could disentangle this conglomerate and represent specific ties more explicitly to focus on a particular connection layer, e.g., connections between co-workers implying network communities that represent job categories. Early contributions to development economics on analogous 'demonstration effects' have emphasised the role of international emulation of consumption patterns (Nurkse, 1953), while our perspective is much more granular and focuses on individuals. Exploring the interplay of aggregate frames of reference between countries and emulation of individual consumption within them might prove to be instructive to investigate cross-country convergence to account for the possibility that the two reinforce each other.

Furthermore, since the model is static, time does not play any role within the model per construction. Simon's (1962) classic cautions us that the dynamic adjustment of different complex systems in response to exogenous disturbances might also operate on vastly different time scales, dependent on its decomposability properties. In this sense, the persistence of the effect of inequality on savings is a function of the time difference between consumption adjustment and network reshuffling. A possible remedy for this could be to explicitly model both consumption decisions, constrained by the availability of credit, and homophilic choice, constrained by mobility, within a full-fledged macro model. One promising avenue for empirical determination of the relevant time scales could be exploiting the time-variation in urbanisation rates. Veblen conjectured already in 1899 that cities are a powerful tool for diversifying perceptions, since within them, "the mobility of the population is greatest" (Veblen, 1899, p. 66). The yet relatively scarce empirical evidence for Vietnam and Central and Eastern Europe suggests that perceived income diversity is indeed increasing with urbanisation (Mahajan et al., 2014; Binelli and Loveless, 2016). Apart from theoretical models, we hope to invite empirical studies, especially at the intersection of urban or spatial economics, inequality research and consumer theory.

References

S. Agarwal, V. Mikhed, and B. Scholnick. Peers' income and financial distress: Evidence from lottery winners and neighboring bankruptcies. *Review of Financial Studies*, 33(1): 433–472, 2020.

G. A. Akerlof. Social distance and social decisions. *Econometrica*, 65(5):1005–1027, 1997.

K. Annen. Social capital, inclusive networks, and economic performance. *Journal of Economic Behavior & Organization*, 50(4):449–463, 2003.

M. Anufriev, A. Chernulich, and J. Tuinstra. A laboratory experiment on the heuristic switching model. *Journal of Economic Dynamics and Control*, 91:21–42, 2018.

F. Ascione and M. Schnetzer. Out of balance? Revisiting the nexus of income inequality, household debt and current account imbalances after the Great Recession. *Cambridge Journal of Economics*, 46:133–165, 2022.

E. Battistin, R. Blundell, and A. Lewbel. Why is consumption more log normal than income? Gibrat's law revisited. *Journal of Political Economy*, 117(6):1140–1154, 2009.

J. Behringer and T. Van Treeck. Income distribution and the current account. *Journal of International Economics*, 114:238–254, 2018.

J. Behringer and T. van Treeck. Income distribution and growth models: a sectoral balances approach. *Politics & Society*, 47(3):303–332, 2019.

J. Behringer and T. van Treeck. Varieties of capitalism and growth regimes: the role of income distribution. *Socio-Economic Review*, 20(3):1249–1286, 2022.

J. Behringer, M. Gonzalez-Granda, and T. Van Treeck. Varieties of the rat race: working hours in the age of abundance. *Socio-Economic Review*, in print, 2023.

M. Bertrand and A. Morse. Trickle-down consumption. *Review of Economics and Statistics*, 98(5):863–879, 2016.

C. Binelli and M. Loveless. The urban–rural divide: Perceptions of income and social inequality in Central and Eastern Europe. *Economics of Transition*, 24(2):211–231, 2016.

P. Bofinger and P. Scheuermeyer. Income distribution and aggregate saving: A non-monotonic relationship. *Review of Income and Wealth*, 65(4):872–907, 2019.

J. Bricker, J. Krimmel, and R. Ramcharan. Signaling status: The impact of relative income on household consumption and financial decisions. *Management Science*, Articles in Advance:1–17, 2020.

M. Brzozowski, M. Gervais, P. Klein, and M. Suzuki. Consumption, income, and wealth inequality in Canada. *Review of Economic Dynamics*, 13(1):52–75, 2010.

M. D. Carr and A. Jayadev. Relative income and indebtedness: Evidence from panel data. *Review of Income and Wealth*, 61(4):759–772, 2015.

C. D. Carroll, J. Overland, and D. N. Weil. Comparison utility in a growth model. *Journal of Economic Growth*, 2(4):339–367, 1997.

D. Cepić and Ž. Tonković. How social ties transcend class boundaries? Network variability as tool for exploring occupational homophily. *Social Networks*, 62:33–42, 2020.

A. S. Chakrabarti, A. Chatterjee, T. Nandi, A. Ghosh, and A. Chakraborti. Quantifying invariant features of within-group inequality in consumption across groups. *Journal of Economic Interaction and Coordination*, 13(3):469–490, 2018.

A. Chandra. Labor-market dropouts and the racial wage gap: 1940–1990. *American Economic Review*, 90(2):333–338, 2000.

W.-H. Chen, J. Myles, and G. Picot. Why have poorer neighbourhoods stagnated economically while the richer have flourished? Neighbourhood income inequality in canadian cities. *Urban Studies*, 49(4):877–896, 2012.

R. Chetty, M. O. Jackson, T. Kuchler, J. Stroebel, N. Hendren, R. B. Fluegge, S. Gong, F. Gonzalez, A. Grondin, M. Jacob, et al. Social capital II: determinants of economic connectedness. *Nature*, 608:122–134, 2022.

G. Choi. Revisiting the redistribution hypothesis with perceived inequality and redistributive preferences. *European Journal of Political Economy*, 58:220–244, 2019. ISSN 01762680.

F. Clementi and L. Gianmoena. Modeling the joint distribution of income and consumption in Italy: A copula-based approach with κ-generalized margins. In M. Gallegati, A. Palestrini, and A. Russo, editors, *Introduction to Agent-Based Economics*, pages 191–228. Elsevier, 2017.

J. C. Cuaresma, J. Kubala, and K. Petrikova. Does income inequality affect aggregate consumption? Revisiting the evidence. *Empirical Economics*, 55(2):905–912, 2018.

G. De Giorgi, A. Frederiksen, and L. Pistaferri. Consumption network effects. *Review of Economic Studies*, 87(1):130–163, 2020.

A. De Paula. Econometrics of network models. In *Advances in Economics and Econometrics: Theory and Applications: Eleventh World Congress*, volume 1, pages 268–323. Cambridge University Press, 2017.

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

M. Drechsel-Grau and K. D. Schmid. Consumption–savings decisions under upward-looking comparisons. *Journal of Economic Behavior & Organization*, 106:254–268, 2014.

J. S. Duesenberry. *Income, saving, and the theory of consumer behavior*. Harvard University Press, 1949.

K. E. Dynan, J. Skinner, and S. P. Zeldes. Do the rich save more? *Journal of Political Economy*, 112(2):397–444, 2004.

G. Fagiolo, L. Alessi, M. Barigozzi, and M. Capasso. On the distributional properties of household consumption expenditures: The case of Italy. *Empirical Economics*, 38(3): 717–741, 2010.

G. Fagiolo, M. Guerini, F. Lamperti, A. Moneta, and A. Roventini. Validation of agent-based models in economics and finance. In C. Beisbart and N. J. Saam, editors, *Computer Simulation Validation: Fundamental Concepts, Methodological Frameworks, and Philosophical Perspectives*, pages 763–787. Springer International Publishing, Cham, 2019.

A. Falk and M. Knell. Choosing the Joneses: Endogenous goals and reference standards. *Scandinavian Journal of Economics*, 106(3):417–435, 2004.

R. H. Frank. *Choosing the right pond: Human behavior and the quest for status*. Oxford University Press, 1985.

R. H. Frank, A. S. Levine, and O. Dijk. Expenditure cascades. *Review of Behavioral Economics*, 1(1–2):55–73, 2014. ISSN 2326-6198.

R. Franke and F. Westerhoff. Taking stock: A rigorous modelling of animal spirits in macroeconomics. *Journal of Economic Surveys*, 31(5):1152–1182, 2017.

M. Friedman. *Theory of the consumption function*. Princeton University Press, 1957.

C. Gräbner. How to relate models to reality? An epistemological framework for the validation and verification of computational models. *Journal of Artificial Societies and Social Simulation*, 21(3), 2018.

T. Grüne-Yanoff and P. Verreault-Julien. How-possibly explanations in economics: anything goes? *Journal of Economic Methodology*, 28(1):114–123, 2021.

X. Gu and B. Huang. Does inequality lead to a financial crisis? Revisited. *Review of Development Economics*, 18(3):502–516, 2014.

X. Gu, B. Dong, and B. Huang. Inequality, saving and global imbalances: A new theory with evidence from OECD and Asian countries. *The World Economy*, 38(1):110–135, 2015.

R. E. Hall. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987, 1978.

P. Harting and D. Radi. Residential segregation: The role of inequality and housing subsidies. *Journal of Economic Behavior & Organization*, 178:801–819, 2020.

O. Heffetz. A test of conspicuous consumption: Visibility and income elasticities. *Review of Economics and Statistics*, 93(4):1101–1117, 2011.

R. Hegselmann. Thomas c. schelling and james m. sakoda: The intellectual, technical, and social history of a model. *Journal of Artificial Societies and Social Simulation*, 20(3), 2017. doi: 10.18564/jasss.3511.

D. F. Hendry and G. E. Mizon. Unpredictability in economic analysis, econometric modeling and forecasting. *Journal of Econometrics*, 182(1):186–195, 2014.

F. Hirsch. *Social limits to growth*. Harvard University Press, 1976.

S. D. Hoffman and G. J. Duncan. Multinomial and conditional logit discrete-choice models in demography. *Demography*, 25(3):415–427, 1988.

M. Hohnisch, S. Pittnauer, and M. Chakrabarty. Empirical regularities in distributions of individual consumption expenditure. *International Journal of Modern Physics C*, 13 (04):541–549, 2002.

R. B. Howarth. Status effects and environmental externalities. *Ecological Economics*, 16 (1):25–34, 1996.

T. Jappelli and L. Pistaferri. Does consumption inequality track income inequality in Italy? *Review of Economic Dynamics*, 13(1):133–153, 2010.

T. Jappelli and L. Pistaferri. Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–36, 2014.

D. Jinkins. Conspicuous consumption in the United States and China. *Journal of Economic Behavior & Organization*, 127:115–132, 2016.

N. Karlsson, T. Gärling, P. Dellgran, and B. Klingander. Social comparison and consumer behavior: When feeling richer or poorer than others is more important than being so. *Journal of Applied Social Psychology*, 35(6):1206–1222, 2005.

J. M. Keynes. *Essays in Persuasion: The Collected Writings, vol. 9 London*. Macmillan, 1972.

A. P. Kirman. Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136, 1992.

M. Klein. Inequality and household debt: a panel cointegration analysis. *Empirica*, 42(2): 391–412, 2015.

J. Koo and Y. Song. The relationship between income inequality and aggregate saving: an empirical analysis using cross-country panel data. *Applied Economics*, 48(10):892–901, 2016.

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

D. Krueger and F. Perri. Does income inequality lead to consumption inequality? Evidence and theory. *The Review of Economic Studies*, 73(1):163–193, 2006.

S. Kuznets. *Uses of National Income in Peace and War*. National Bureau of Economic Research, New York, 1942.

P. F. Lazarsfeld, R. K. Merton, et al. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society*, 18(1):18–66, 1954.

A. Leigh and A. Posso. Top incomes and national savings. *Review of Income and Wealth*, 55(1):57–74, 2009.

Y. Leo, E. Fleury, J. I. Alvarez-Hamelin, C. Sarraute, and M. Karsai. Socioeconomic correlations and stratification in social-communication networks. *Journal of The Royal Society Interface*, 13(125):20160598, dec 2016.

K.-H. Li and C. T. Li. Linear combination of independent exponential random variables. *Methodology and Computing in Applied Probability*, 21(1):253–277, 2019.

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

S. Mahajan, V. T. Dinh, H. Rab, V. Q. Trieu, L. V. Nguyen, V. Kwakwa, and M. Verghis. Taking stock: an update on Vietnam's recent economic developments. Technical report, Washington, D.C.: World Bank Group, 2014.

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

T. Malacarne. Rich friends, poor friends: Inter–socioeconomic status friendships in secondary school. *Socius*, 3:2378023117736994, 2017.

C. F. Manski and D. McFadden. *Structural analysis of discrete data with econometric applications*. MIT press Cambridge, MA, 1981.

R. Mason. The social significance of consumption: James Duesenberry's contribution to consumer theory. *Journal of Economic Issues*, 34(3):553–572, 2000.

A. Mayer and S. L. Puller. The old boy (and girl) network: Social network formation on university campuses. *Journal of Public Economics*, 92(1-2):329–347, 2008.

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

D. M. Mayerhoffer and J. Schulz. Marginalisation and misperception: Perceiving gender and racial wage gaps in ego networks. In R. M. Benito, C. Cherifi, H. Cherifi, E. Moro, L. M. Rocha, and M. Sales-Pardo, editors, *Complex Networks & Their Applications X*, pages 768–779, Cham, 2022b. Springer International Publishing. ISBN 978-3-030-93409-5.

D. M. Mayerhoffer and J. Schulz. Inequality-Perception, 1 2023. URL `https://github.com/mayerhoffer/Inequality-Perception`.

J. M. McPherson and L. Smith-Lovin. Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *American Sociological Review*, 52 (3):370–379, 1987.

C. Meghir. A retrospective on Friedman's theory of permanent income. *Economic Journal*, 114(496):F293–F306, 2004.

J.-O. Menz. Uncertainty, social norms and consumption theory: Post and New Keynesian approaches. *European Journal of Economics and Economic Policies: Intervention*, 7(1): 125–146, 2010.

M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet mathematics*, 1(2):226–251, 2004.

R. Nurkse. *Problems of capital formation in underdeveloped countries*. Oxford University Press, 1953.

T. I. Palley. The relative permanent income theory of consumption: A synthetic Keynes–Duesenberry–Friedman model. *Review of Political Economy*, 22(1):41–56, 2010.

M. H. Pesaran. *Time series and panel data econometrics*. Oxford University Press, 2015.

M. H. Pesaran and M. Weale. Survey expectations. In C. W. J. Granger, G. Elliott, and A. Timmermann, editors, *Handbook of economic forecasting*, volume 1, pages 715–776. Elsevier, 2006.

L. A. Petach and D. Tavani. Consumption externalities and growth: Theory and evidence for the United States. *Journal of Economic Behavior & Organization*, 183:976 – 997, 2021.

W. V. Quine. On empirically equivalent systems of the world. *Erkenntnis*, 9(3):313–328, 1975.

C. Quintana-Domeque and J. Wohlfart. "Relative concerns for consumption at the top": An intertemporal analysis for the UK. *Journal of Economic Behavior & Organization*, 129:172–194, 2016.

J. M. Sakoda. Minidoka: An analysis of changing patterns of social interaction. University of California, Berkeley, 1949.

J. M. Sakoda. The checkerboard model of social interaction. *The Journal of Mathematical Sociology*, 1(1):119–132, jan 1971. doi: 10.1080/0022250x.1971.9989791.

T. C. Schelling. Dynamic models of segregation. *Journal of mathematical sociology*, 1(2): 143–186, 1971.

J. Schulz and M. Milaković. How wealthy are the rich? *Review of Income and Wealth*, 69 (1):100–123, 2023.

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

A. Shaikh, J. E. Jacobo, et al. Economic arbitrage and the econophysics of income inequality. *Review of Behavioral Economics*, 7:299–315, 2020.

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

H. Simon. The architecture of complexity. *Proceedings of American Philosophical Association*, 106(6), 1962.

A. Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Republished (1991), Everyman's Library Classics Series, London: Home University Library, 1776.

S. J. Solnick and D. Hemenway. Are positional concerns stronger in some domains than in others? *American Economic Review*, 95(2):147–151, 2005.

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

R. J. Thomas. Sources of friendship and structurally induced homophily across the life course. *Sociological Perspectives*, 62(6):822–843, 2019.

G. Tóth, J. Wachs, R. Di Clemente, Á. Jakobi, B. Ságvári, J. Kertész, and B. Lengyel. Inequality is rising where social network segregation interacts with urban topology. *Nature Communications*, 12(1):1–9, 2021.

T. Van Treeck. Did inequality cause the US financial crisis? *Journal of Economic Surveys*, 28(3):421–448, 2014.

T. Veblen. *The Theory of the Leisure Class*. Republished (2001), New York: The Modern Library, 1899.

M. Virtanen, M. Kivimäki, A. Kouvonen, M. Elovainio, A. Linna, T. Oksanen, and J. Vahtera. Average household income, crime, and smoking behaviour in a local area: The finnish 10-town study. *Social Science & Medicine*, 64(9):1904–1913, may 2007.

K. A. Weeden and B. Cornwell. The small-world network of college classes: implications for epidemic spread on a university campus. *Sociological Science*, 7:222–241, 2020.

K. Wilcox and A. T. Stephen. Are close friends the enemy? Online social networks, self-esteem, and self-control. *Journal of Consumer Research*, 40(1):90–103, 2013.

R. Wildauer and E. Stockhammer. Expenditure cascades, low interest rates, credit deregulation or property booms? Determinants of household debt in OECD countries. *Review of Behavioral Economics*, 5(2):85–121, 2018.

U. Witt. Learning to consume–a theory of wants and the growth of demand. *Journal of evolutionary economics*, 11(1):23–36, 2001.

W.-X. Zhou, D. Sornette, R. A. Hill, and R. I. Dunbar. Discrete hierarchical organization of social group sizes. *Proceedings of the Royal Society B: Biological Sciences*, 272(1561): 439–444, 2005.